

# EXHIBITS A1-A6

## (Part 10 of 13)

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record				
<p><b>distance (EIGRP)</b></p> <p>To allow the use of two administrative distances—internal and external—for the Enhanced Interior Gateway Routing Protocol (EIGRP) that could provide a better route to a node, use the <code>distance</code> command. To reset to default, use the <code>no</code> form of this command.</p> <pre>distance internal-distance external-distance no distance</pre> <p><b>Syntax Description</b></p> <table border="0"> <tr> <td style="vertical-align: top;"><code>internal-distance</code></td> <td>Administrative distance for EIGRP internal routes. Internal routes are routes that are learned from another entity within the same autonomous system (AS). The distance can be a value from 1 to 255. The default value is 90.</td> </tr> <tr> <td style="vertical-align: top;"><code>external-distance</code></td> <td>Administrative distance for EIGRP external routes. External routes are routes for which the best path is learned from a source external to this autonomous system. The distance can be a value from 1 to 255. The default value is 170.</td> </tr> </table> <p><b>Defaults</b></p> <pre>internal-distance: 90 external-distance: 170</pre> <p>Cisco IOS IP Routing: EIGRP Command Reference (2009), at IRE-33.</p>	<code>internal-distance</code>	Administrative distance for EIGRP internal routes. Internal routes are routes that are learned from another entity within the same autonomous system (AS). The distance can be a value from 1 to 255. The default value is 90.	<code>external-distance</code>	Administrative distance for EIGRP external routes. External routes are routes for which the best path is learned from a source external to this autonomous system. The distance can be a value from 1 to 255. The default value is 170.	<p><b>distance bgp</b></p> <p>The <code>distance bgp</code> command assigns an administrative distance to routes that the switch learns through BGP. Routers use administrative distances to select a route when two protocols provide routing information to the same destination. Distance values range from 1 to 255; lower distance values correspond to higher reliability. BGP routing tables do not include routes with a distance of 255.</p> <p>The <code>distance</code> command assigns distance values to external, internal, and local BGP routes:</p> <ul style="list-style-type: none"> <li><b>external:</b> External routes are routes for which the best path is learned from a neighbor external to the autonomous system. Default distance is 200.</li> <li><b>internal:</b> Internal routes are routes learned from a BGP entity within the same autonomous system. Default distance is 200.</li> <li><b>local:</b> Local routes are networks listed with a network router configuration command for that router or for networks that are redistributed from another process. Default distance is 200.</li> </ul> <p>The <code>no distance bgp</code> and <code>default distance bgp</code> commands restore the default administrative distances by removing the <code>distance bgp</code> command from <i>running-config</i>.</p> <pre>Platform      all Command Mode  Router-BGP Configuration</pre> <p><b>Command Syntax</b></p> <pre>distance bgp external_dist [INTERNAL_LOCAL] no distance bgp default distance bgp</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li><b>external_dist</b> distance assigned to external routes. Values range from 1 to 255.</li> <li><b>INTERNAL_LOCAL</b> distance assigned to internal and local routes. Values for both routes range from 1 to 255. Options include: <ul style="list-style-type: none"> <li>&lt;no parameter&gt; <code>external_dist</code> value is assigned to internal and local routes.</li> <li><code>internal_dist local_dist</code> values assigned to internal (<code>internal_dist</code>) and local (<code>local_dist</code>) routes.</li> </ul> </li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1583.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1360; Arista User Manual, v. 4.11.1 (1/11/13), at 1106; Arista User Manual v. 4.10.3 (10/22/12), at 918; Arista User Manual v. 4.9.3.2 (5/3/12), at 684; Arista User Manual v. 4.8.2 (11/18/11), at 514; Arista User Manual v. 4.7.3 (7/18/11), at 379.</p>	Dkt. 419-10 at PDF p. 319
<code>internal-distance</code>	Administrative distance for EIGRP internal routes. Internal routes are routes that are learned from another entity within the same autonomous system (AS). The distance can be a value from 1 to 255. The default value is 90.					
<code>external-distance</code>	Administrative distance for EIGRP external routes. External routes are routes for which the best path is learned from a source external to this autonomous system. The distance can be a value from 1 to 255. The default value is 170.					

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
<p>When you configure the ip command on an interface, the handling of proxy Address Resolution Protocol (ARP) requests changes (unless proxy ARP was disabled). Hosts send ARP requests to map an IP address to a MAC address. The GLBP gateway intercepts the ARP requests and replies to the ARP requests on behalf of the connected nodes. If a forwarder in the GLBP group is active, proxy ARP requests are answered using the MAC address of the first active forwarder in the group. If no forwarder is active, proxy ARP responses are suppressed.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 256.</p>	<p><b>Displaying ARP Entries</b></p> <p>The <code>show ip arp</code> command displays ARP cache entries that map an IP address to a corresponding MAC address. The table displays addresses by their host names when the command includes the <code>resolve</code> argument.</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1225.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1038; Arista User Manual, v. 4.11.1 (1/11/13), at 840; Arista User Manual v. 4.10.3 (10/22/12), at 687.</p>	Dkt. 419-10 at PDF p. 320
<p>When you configure the ip command on an interface, the handling of proxy Address Resolution Protocol (ARP) requests changes (unless proxy ARP was disabled). Hosts send ARP requests to map an IP address to a MAC address. The GLBP gateway intercepts the ARP requests and replies to the ARP requests on behalf of the connected nodes. If a forwarder in the GLBP group is active, proxy ARP requests are answered using the MAC address of the first active forwarder in the group. If no forwarder is active, proxy ARP responses are suppressed.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 5.x (2010), at L3-236.</p>	<p><b>Displaying ARP Entries</b></p> <p>The <code>show ip arp</code> command displays ARP cache entries that map an IP address to a corresponding MAC address. The table displays addresses by their host names when the command includes the <code>resolve</code> argument.</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1225.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1038; Arista User Manual, v. 4.11.1 (1/11/13), at 840; Arista User Manual v. 4.10.3 (10/22/12), at 687.</p>	Dkt. 419-10 at PDF p. 320
<p>When you configure the ip command on an interface, the handling of proxy Address Resolution Protocol (ARP) requests changes (unless proxy ARP was disabled). Hosts send ARP requests to map an IP address to a MAC address. The GLBP gateway intercepts the ARP requests and replies to the ARP requests on behalf of the connected nodes. If a forwarder in the GLBP group is active, proxy ARP requests are answered using the MAC address of the first active forwarder in the group. If no forwarder is active, proxy ARP responses are suppressed.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 4.x (2008), at L3-143.</p>	<p><b>Displaying ARP Entries</b></p> <p>The <code>show ip arp</code> command displays ARP cache entries that map an IP address to a corresponding MAC address. The table displays addresses by their host names when the command includes the <code>resolve</code> argument.</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1225.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1038; Arista User Manual, v. 4.11.1 (1/11/13), at 840; Arista User Manual v. 4.10.3 (10/22/12), at 687.</p>	Dkt. 419-10 at PDF p. 320

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
<p>Address Resolution Protocol (ARP) is an Internet protocol used to map an IP address to a MAC address. ARP finds the MAC address, also known as the hardware address, of an IP-routed host from its known IP address and maintains this mapping information in a table. The router uses this IP address and MAC address mapping information to send IP packets to the next-hop router in the network.</p> <p>Cisco IOS IP Addressing Services Configuration Guide (2009), at CS1-CL1-00061623.</p>	<p><b>Displaying ARP Entries</b></p> <p>The <code>show ip arp</code> command displays ARP cache entries that map an IP address to a corresponding MAC address. The table displays addresses by their host names when the command includes the <code>resolve</code> argument.</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1225.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1038; Arista User Manual, v. 4.11.1 (1/11/13), at 840; Arista User Manual v. 4.10.3 (10/22/12), at 687.</p>	Dkt. 419-10 at PDF p. 321
<p><b>Expanded Community Lists</b></p> <p>Expanded community lists are used to filter communities using a regular expression. Regular expressions are used to configure patterns to match community attributes. The order for matching using the * or + character is longest construct first. Nested constructs are matched from the outside in. Concatenated constructs are matched beginning at the left side. If a regular expression can match two different parts of an input string, it will match the earliest part first.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 274.</p>	<p>The order for matching using the * or + character is longest construct first. Nested constructs are matched from the outside in. Concatenated constructs are matched beginning at the left side. If a regular expression can match two different parts of an input string, it matches the earliest part first.</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 107.</p> <p><i>See also</i> Arista User Manual v. 4.13.6F (4/14/2014), at 105; Arista User Manual, v. 4.11.1 (1/11/13), at 65; Arista User Manual v. 4.12.3 (7/17/13), at 95; Arista User Manual v. 4.10.3 (10/22/12), at 57; Arista User Manual v. 4.9.3.2 (5/3/12), at 53; Arista User Manual v. 4.8.2 (11/18/11), at 49.</p>	Dkt. 419-10 at PDF p. 321
<p><b>Expanded Community Lists</b></p> <p>Expanded community lists are used to filter communities using a regular expression. Regular expressions are used to configure patterns to match community attributes. The order for matching using the * or + character is longest construct first. Nested constructs are matched from the outside in. Concatenated constructs are matched beginning at the left side. If a regular expression can match two different parts of an input string, it will match the earliest part first.</p> <p>Cisco IOS IP Routing: BGP Command Reference, (2009), at 274.</p>	<p>The order for matching using the * or + character is longest construct first. Nested constructs are matched from the outside in. Concatenated constructs are matched beginning at the left side. If a regular expression can match two different parts of an input string, it matches the earliest part first.</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 107.</p> <p><i>See also</i> Arista User Manual v. 4.13.6F (4/14/2014), at 105; Arista User Manual, v. 4.11.1 (1/11/13), at 65; Arista User Manual v. 4.12.3 (7/17/13), at 95; Arista User Manual v. 4.10.3 (10/22/12), at 57; Arista User Manual v. 4.9.3.2 (5/3/12), at 53; Arista User Manual v. 4.8.2 (11/18/11), at 49.</p>	Dkt. 419-10 at PDF p. 321

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
<p>Static routes have a default administrative distance of 1. If you want a dynamic routing protocol to take precedence over a static route, you must configure the static route preference argument to be greater than the administrative distance of the dynamic routing protocol. For example, routes derived with Enhanced Interior Gateway Routing Protocol (EIGRP) have a default administrative distance of 100. To have a static route that would be overridden by an EIGRP dynamic route, specify an administrative distance greater than 100.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 337.</p>	<p>Static routes have a default administrative distance of 1. Assigning a higher administrative distance to a static route configures it to be overridden by dynamic routing data. For example, a static route with a distance value of 200 is overridden by OSPF intra-area routes with a default distance of 110.</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1226.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1082; Arista User Manual, v. 4.11.1 (1/11/13), at 860; Arista User Manual v. 4.10.3 (10/22/12), at 683.</p>	Dkt. 419-10 at PDF p. 322
<p>Static routes have a default administrative distance of 1. If you want a dynamic routing protocol to take precedence over a static route, you must configure the static route preference argument to be greater than the administrative distance of the dynamic routing protocol. For example, routes derived with Enhanced Interior Gateway Routing Protocol (EIGRP) have a default administrative distance of 100. To have a static route that would be overridden by an EIGRP dynamic route, specify an administrative distance greater than 100.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 5.x (2010), at L3-311.</p>	<p>Static routes have a default administrative distance of 1. Assigning a higher administrative distance to a static route configures it to be overridden by dynamic routing data. For example, a static route with a distance value of 200 is overridden by OSPF intra-area routes with a default distance of 110.</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1226.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1082; Arista User Manual, v. 4.11.1 (1/11/13), at 860; Arista User Manual v. 4.10.3 (10/22/12), at 683.</p>	Dkt. 419-10 at PDF p. 322

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
<p><b>is-type</b></p> <p>To configure the routing level for an instance of the Intermediate System-to-Intermediate System (IS-IS) routing process, use the <b>is-type</b> command. To reset the default value, use the <b>no</b> form of this command.</p> <pre>is-type {level-1   level-1-2   level-2} no is-type {level-1   level-1-2   level-2}</pre> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 407.</p>	<p><b>is-type</b></p> <p>The <b>is-type</b> command configures the routing level for an instance of the IS-IS routing instance.</p> <p>Platform all Command Mode Router-IS-IS Configuration</p> <p><b>Command Syntax</b></p> <pre>is-type LAYER_VALUE</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>LAYER_VALUE</b> layer value. Options include: <ul style="list-style-type: none"> <li>— <b>level-1</b> The switch operates as a <b>Level-1</b> (intra-area) router.</li> <li>— <b>level-2</b> The switch operates as a <b>Level-2</b> (inter-area) router.</li> </ul> </li> </ul> <p><b>Example</b></p> <ul style="list-style-type: none"> <li>• These commands configure Level 2 routing on interface Ethernet 5.</li> </ul> <pre>switch(config)#router isis Osiris switch(config-router-isis)#is-type level-2 switch(config-router-isis)# </pre> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1691.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1451.</p>	Dkt. 419-10 at PDF p. 323

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
<p><b>is-type</b></p> <p>To configure the routing level for an instance of the Intermediate System-to-Intermediate System (IS-IS) routing process, use the <b>is-type</b> command. To reset the default value, use the <b>no</b> form of this command.</p> <pre>is-type {level-1   level-1-2   level-2} no is-type {level-1   level-1-2   level-2}</pre> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 5.x (2010), at L3-373.</p>	<p><b>is-type</b></p> <p>The <b>is-type</b> command configures the routing level for an instance of the IS-IS routing instance.</p> <p>Platform all Command Mode Router-IS-IS Configuration</p> <p><b>Command Syntax</b></p> <pre>is-type LAYER_VALUE</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>LAYER_VALUE</b> layer value. Options include: <ul style="list-style-type: none"> <li>— <b>level-1</b> The switch operates as a <b>Level-1</b> (intra-area) router.</li> <li>— <b>level-2</b> The switch operates as a <b>Level-2</b> (inter-area) router.</li> </ul> </li> </ul> <p><b>Example</b></p> <ul style="list-style-type: none"> <li>• These commands configure Level 2 routing on interface Ethernet 5.</li> </ul> <pre>switch(config)#router isis Osiris switch(config-router-isis)#is-type level-2 switch(config-router-isis)# </pre> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1691.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1451.</p>	Dkt. 419-10 at PDF p. 324

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
<p><b>is-type</b></p> <p>To configure the routing level for an instance of the Intermediate System-to-Intermediate System (IS-IS) routing process, use the <b>is-type</b> command. To reset the default value, use the <b>no</b> form of this command.</p> <pre>is-type {level-1   level-1-2   level-2} no is-type {level-1   level-1-2   level-2}</pre> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 4.x (2008), at L3-208.</p>	<p><b>is-type</b></p> <p>The <b>is-type</b> command configures the routing level for an instance of the IS-IS routing instance.</p> <p>Platform all Command Mode Router-IS-IS Configuration</p> <p><b>Command Syntax</b></p> <pre>is-type LAYER_VALUE</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>LAYER_VALUE</b> layer value. Options include: <ul style="list-style-type: none"> <li>— <b>level-1</b> The switch operates as a <b>Level-1</b> (intra-area) router.</li> <li>— <b>level-2</b> The switch operates as a <b>Level-2</b> (inter-area) router.</li> </ul> </li> </ul> <p><b>Example</b></p> <ul style="list-style-type: none"> <li>• These commands configure Level 2 routing on interface Ethernet 5.</li> </ul> <pre>switch(config)#router isis Osiris switch(config-router-isis)#is-type level-2 switch(config-router-isis)# </pre> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1691.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1451.</p>	Dkt. 419-10 at PDF p. 325

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record				
<p><b>is-type</b></p> <div style="border: 1px solid red; padding: 5px;"> <p>To configure the routing level for an instance of the Intermediate System-to-Intermediate System (IS-IS) routing process, use the <b>is-type</b> command in router configuration mode. To reset the default value, use the <b>no</b> form of this command.</p> <pre>is-type [level-1   level-1-2   level-2-only] no is-type [level-1   level-1-2   level-2-only]</pre> </div> <p>Cisco IOS IP Routing: ISIS Command Reference (2009), at IRS-73.</p>	<p><b>is-type</b></p> <p>The <b>is-type</b> command configures the routing level for an instance of the IS-IS routing instance.</p> <table style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 15%;">Platform</td> <td style="width: 15%;">all</td> </tr> <tr> <td>Command Mode</td> <td>Router-ISIS Configuration</td> </tr> </table> <p><b>Command Syntax</b></p> <pre>is-type LAYER_VALUE</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>LAYER_VALUE</b> layer value. Options include:       <ul style="list-style-type: none"> <li>— <b>level-1</b> The switch operates as a <b>Level-1</b> (intra-area) router.</li> <li>— <b>level-2</b> The switch operates as a <b>Level-2</b> (inter-area) router.</li> </ul> </li> </ul> <p><b>Example</b></p> <ul style="list-style-type: none"> <li>• These commands configure Level 2 routing on interface Ethernet 5.</li> </ul> <pre>switch(config)#router isis Osiris switch(config-router-isis)#is-type level-2 switch(config-router-isis)# </pre> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1691.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1451.</p>	Platform	all	Command Mode	Router-ISIS Configuration	Dkt. 419-10 at PDF p. 326
Platform	all					
Command Mode	Router-ISIS Configuration					

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record						
<p><b>isis hello-multiplier</b></p> <p>To specify the number of Intermediate System-to-Intermediate System (IS-IS) hello packets a neighbor must miss before the router should declare the adjacency as down, use the <b>isis hello-multiplier</b> command. To restore the default value, use the no form of this command.</p> <pre>isis hello-multiplier multiplier [level-1   level-2] no isis hello-multiplier [level-1   level-2]</pre> <table border="1"> <tr> <td><b>Syntax Description</b></td> <td> <b>multiplier</b> Integer value. Range: 3 to 1000. Default: 3.  <b>level-1</b> Configures the hello multiplier independently for Level 1 adjacencies.  <b>level-2</b> Configures the hello multiplier independently for Level 2 adjacencies.         </td> </tr> </table> <p><b>Command Default</b></p> <p>The default settings are as follows:</p> <ul style="list-style-type: none"> <li>• <b>multiplier:</b> 3</li> <li>• Level 1 and Level 2</li> </ul> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 423.</p>	<b>Syntax Description</b>	<b>multiplier</b> Integer value. Range: 3 to 1000. Default: 3. <b>level-1</b> Configures the hello multiplier independently for Level 1 adjacencies. <b>level-2</b> Configures the hello multiplier independently for Level 2 adjacencies.	<p><b>isis hello-multiplier</b></p> <p>The <b>isis hello-multiplier</b> command specifies the number of IS-IS hello packets a neighbor must miss before the device should declare the adjacency as down.</p> <p>Each hello packet contains a hold time. The hold time informs the receiving devices how long to wait without seeing another hello from the sending device before considering the sending device down. The <b>isis hello-multiplier</b> command is used to calculate the hold time announced in hello packets by multiplying this number with the configured <b>isis hello-interval</b>.</p> <p>The no <b>isis hello-multiplier</b> and default <b>isis hello-multiplier</b> commands restore the default hello interval of 3 on the configuration mode interface by removing the <b>isis hello-multiplier</b> command from <i>running-config</i>.</p> <table> <tr> <td><b>Platform</b></td> <td>all</td> </tr> <tr> <td><b>Command Mode</b></td> <td>Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration</td> </tr> </table> <p><b>Command Syntax</b></p> <pre>isis hello-multiplier factor no isis hello-multiplier default isis hello-multiplier</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>factor</b> hello multiplier. Values range from 3 to 100; default is 3</li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1685.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1447.</p>	<b>Platform</b>	all	<b>Command Mode</b>	Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration	Dkt. 419-10 at PDF p. 327
<b>Syntax Description</b>	<b>multiplier</b> Integer value. Range: 3 to 1000. Default: 3. <b>level-1</b> Configures the hello multiplier independently for Level 1 adjacencies. <b>level-2</b> Configures the hello multiplier independently for Level 2 adjacencies.							
<b>Platform</b>	all							
<b>Command Mode</b>	Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration							

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record						
<p><b>isis hello-multiplier</b></p> <p>To specify the number of Intermediate System-to-Intermediate System (IS-IS) hello packets a neighbor must miss before the router should declare the adjacency as down, use the <b>isis hello-multiplier</b> command. To restore the default value, use the no form of this command.</p> <pre>isis hello-multiplier multiplier [level-1   level-2] no isis hello-multiplier [level-1   level-2]</pre> <table border="1"> <tr> <td><b>Syntax Description</b></td> <td> <b>multiplier</b> Integer value. Range: 3 to 1000. Default: 3.  <b>level-1</b> Configures the hello multiplier independently for Level 1 adjacencies.  <b>level-2</b> Configures the hello multiplier independently for Level 2 adjacencies.         </td> </tr> </table> <p><b>Command Default</b></p> <p>The default settings are as follows:</p> <ul style="list-style-type: none"> <li>• <b>multiplier:</b> 3</li> <li>• Level 1 and Level 2</li> </ul> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 5.x (2010), at L3-389.</p>	<b>Syntax Description</b>	<b>multiplier</b> Integer value. Range: 3 to 1000. Default: 3. <b>level-1</b> Configures the hello multiplier independently for Level 1 adjacencies. <b>level-2</b> Configures the hello multiplier independently for Level 2 adjacencies.	<p><b>isis hello-multiplier</b></p> <p>The <b>isis hello-multiplier</b> command specifies the number of IS-IS hello packets a neighbor must miss before the device should declare the adjacency as down.</p> <p>Each hello packet contains a hold time. The hold time informs the receiving devices how long to wait without seeing another hello from the sending device before considering the sending device down. The <b>isis hello-multiplier</b> command is used to calculate the hold time announced in hello packets by multiplying this number with the configured <b>isis hello-interval</b>.</p> <p>The no <b>isis hello-multiplier</b> and default <b>isis hello-multiplier</b> commands restore the default hello interval of 3 on the configuration mode interface by removing the <b>isis hello-multiplier</b> command from <i>running-config</i>.</p> <table> <tr> <td><b>Platform</b></td> <td>all</td> </tr> <tr> <td><b>Command Mode</b></td> <td>Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration</td> </tr> </table> <p><b>Command Syntax</b></p> <pre>isis hello-multiplier factor no isis hello-multiplier default isis hello-multiplier</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>factor</b> hello multiplier. Values range from 3 to 100; default is 3</li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1685.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1447.</p>	<b>Platform</b>	all	<b>Command Mode</b>	Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration	Dkt. 419-10 at PDF p. 328
<b>Syntax Description</b>	<b>multiplier</b> Integer value. Range: 3 to 1000. Default: 3. <b>level-1</b> Configures the hello multiplier independently for Level 1 adjacencies. <b>level-2</b> Configures the hello multiplier independently for Level 2 adjacencies.							
<b>Platform</b>	all							
<b>Command Mode</b>	Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration							

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record						
<p><b>isis hello-multiplier</b></p> <p>To specify the number of Intermediate System-to-Intermediate System (IS-IS) hello packets a neighbor must miss before the router should declare the adjacency as down, use the <b>isis hello-multiplier</b> command. To restore the default value, use the no form of this command.</p> <pre>isis hello-multiplier multiplier [level-1   level-2] no isis hello-multiplier [level-1   level-2]</pre> <table border="1"> <tr> <td><b>Syntax Description</b></td> <td> <b>multiplier</b> Integer value. Range: 3 to 1000. Default: 3.  <b>level-1</b> Configures the hello multiplier independently for Level 1 adjacencies.  <b>level-2</b> Configures the hello multiplier independently for Level 2 adjacencies.         </td> </tr> </table> <p><b>Command Default</b></p> <p>The default settings are as follows:</p> <ul style="list-style-type: none"> <li>• <b>multiplier:</b> 3</li> <li>• Level 1 and Level 2</li> </ul> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2008), at L3-224.</p>	<b>Syntax Description</b>	<b>multiplier</b> Integer value. Range: 3 to 1000. Default: 3. <b>level-1</b> Configures the hello multiplier independently for Level 1 adjacencies. <b>level-2</b> Configures the hello multiplier independently for Level 2 adjacencies.	<p><b>isis hello-multiplier</b></p> <p>The <b>isis hello-multiplier</b> command specifies the number of IS-IS hello packets a neighbor must miss before the device should declare the adjacency as down.</p> <p>Each hello packet contains a hold time. The hold time informs the receiving devices how long to wait without seeing another hello from the sending device before considering the sending device down. The <b>isis hello-multiplier</b> command is used to calculate the hold time announced in hello packets by multiplying this number with the configured <b>isis hello-interval</b>.</p> <p>The no <b>isis hello-multiplier</b> and default <b>isis hello-multiplier</b> commands restore the default hello interval of 3 on the configuration mode interface by removing the <b>isis hello-multiplier</b> command from <i>running-config</i>.</p> <table> <tr> <td><b>Platform</b></td> <td>all</td> </tr> <tr> <td><b>Command Mode</b></td> <td>Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration</td> </tr> </table> <p><b>Command Syntax</b></p> <pre>isis hello-multiplier factor no isis hello-multiplier default isis hello-multiplier</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>factor</b> hello multiplier. Values range from 3 to 100; default is 3</li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1685.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1447.</p>	<b>Platform</b>	all	<b>Command Mode</b>	Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration	Dkt. 419-10 at PDF p. 329
<b>Syntax Description</b>	<b>multiplier</b> Integer value. Range: 3 to 1000. Default: 3. <b>level-1</b> Configures the hello multiplier independently for Level 1 adjacencies. <b>level-2</b> Configures the hello multiplier independently for Level 2 adjacencies.							
<b>Platform</b>	all							
<b>Command Mode</b>	Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration							

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record						
<p><b>isis hello-multiplier</b></p> <p>To specify the number of Intermediate System-to-Intermediate System (IS-IS) hello packets a neighbor must miss before the router should declare the adjacency as down, use the <b>isis hello-multiplier</b> command. To restore the default value, use the no form of this command.</p> <pre>isis hello-multiplier multiplier {level-1   level-2} no isis hello-multiplier {level-1   level-2}</pre> <table border="1"> <tr> <td><b>Syntax Description</b></td> <td> <b>multiplier</b> Integer value. Range: 3 to 1000. Default: 3.  <b>level-1</b> Configures the hello multiplier independently for Level 1 adjacencies.  <b>level-2</b> Configures the hello multiplier independently for Level 2 adjacencies.         </td> </tr> </table> <p><b>Command Default</b></p> <p>The default settings are as follows:</p> <ul style="list-style-type: none"> <li>• <b>multiplier:</b> 3</li> <li>• Level 1 and Level 2</li> </ul> <p>Cisco IOS IP Routing: ISIS Command Reference (2009), at IRS-54.</p>	<b>Syntax Description</b>	<b>multiplier</b> Integer value. Range: 3 to 1000. Default: 3. <b>level-1</b> Configures the hello multiplier independently for Level 1 adjacencies. <b>level-2</b> Configures the hello multiplier independently for Level 2 adjacencies.	<p><b>isis hello-multiplier</b></p> <p>The <b>isis hello-multiplier</b> command specifies the number of IS-IS hello packets a neighbor must miss before the device should declare the adjacency as down.</p> <p>Each hello packet contains a hold time. The hold time informs the receiving devices how long to wait without seeing another hello from the sending device before considering the sending device down. The <b>isis hello-multiplier</b> command is used to calculate the hold time announced in hello packets by multiplying this number with the configured <b>isis hello-interval</b>.</p> <p>The no <b>isis hello-multiplier</b> and default <b>isis hello-multiplier</b> commands restore the default hello interval of 3 on the configuration mode interface by removing the <b>isis hello-multiplier</b> command from <i>running-config</i>.</p> <table> <tr> <td><b>Platform</b></td> <td>all</td> </tr> <tr> <td><b>Command Mode</b></td> <td>Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration</td> </tr> </table> <p><b>Command Syntax</b></p> <pre>isis hello-multiplier factor no isis hello-multiplier default isis hello-multiplier</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>factor</b> hello multiplier. Values range from 3 to 100; default is 3</li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1685.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1447.</p>	<b>Platform</b>	all	<b>Command Mode</b>	Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration	Dkt. 419-10 at PDF p. 330
<b>Syntax Description</b>	<b>multiplier</b> Integer value. Range: 3 to 1000. Default: 3. <b>level-1</b> Configures the hello multiplier independently for Level 1 adjacencies. <b>level-2</b> Configures the hello multiplier independently for Level 2 adjacencies.							
<b>Platform</b>	all							
<b>Command Mode</b>	Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration							

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record										
<p><b>isis priority</b></p> <p>To configure the priority of designated routers, use the <b>isis priority</b> command in interface configuration mode. To reset the default priority, use the <b>no</b> form of this command.</p> <pre>isis priority number-value [level-1   level-2] no isis priority [level-1   level-2]</pre> <p><b>Syntax Description</b></p> <table border="1"> <tr> <td><b>number-value</b></td> <td>Priority of a router and is a number from 0 to 127. The default value is 64.</td> </tr> <tr> <td><b>level-1</b></td> <td>(Optional) Sets the priority for Level 1 independently.</td> </tr> <tr> <td><b>level-2</b></td> <td>(Optional) Sets the priority for Level 2 independently.</td> </tr> </table> <p><b>Defaults</b></p> <p>Priority of 64 Level 1 and Level 2</p> <p><b>Command Modes</b></p> <p>Interface configuration</p> <p><b>SupportedUserRoles</b></p> <p>network-admin vdc-admin</p> <p><b>Command History</b></p> <table border="1"> <tr> <th>Release</th> <th>Modification</th> </tr> <tr> <td>4.0(1)</td> <td>This command was introduced.</td> </tr> </table> <p><b>Usage Guidelines</b></p> <p>Priorities can be configured for Level 1 and Level 2 independently. Specifying the <b>level-1</b> or <b>level-2</b> keyword resets priority only for Level 1 or Level 2 routing, respectively.</p> <p>The priority is used to determine which router on a LAN will be the designated router or Designated Intermediate System (DIS). The priorities are advertised in the hello packets. The router with the highest priority will become the DIS.</p> <p>In Intermediate System-to-Intermediate System (IS-IS), there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a router with a higher priority comes on line, it will take over the role from the current DIS. In the case of equal priorities, the highest MAC address breaks the tie.</p> <p>This command requires the Enterprise Services license.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 433.</p>	<b>number-value</b>	Priority of a router and is a number from 0 to 127. The default value is 64.	<b>level-1</b>	(Optional) Sets the priority for Level 1 independently.	<b>level-2</b>	(Optional) Sets the priority for Level 2 independently.	Release	Modification	4.0(1)	This command was introduced.	<p><b>isis priority</b></p> <p>The <b>isis priority</b> command configures IS-IS router priority for the configuration mode interface.</p> <p>The priority is used to determine which device will be the Designated Intermediate System (DIS). The device with the highest priority will become the DIS.</p> <p>In IS-IS, there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a device with a higher priority comes on line, it will take over the role from the current DIS.</p> <p>The <b>no isis priority</b> and <b>default isis priority</b> commands restore the default priority (64) on the configuration mode interface.</p> <p><b>Platform</b> all  <b>Command Mode</b> Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration</p> <p><b>Command Syntax</b></p> <pre>isis priority priority_level no isis priority default isis priority</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>priority_level</b> priority level. Value ranges from 0 to 127. Default value is 64.</li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1690.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1450.</p>	Dkt. 419-10 at PDF p. 331
<b>number-value</b>	Priority of a router and is a number from 0 to 127. The default value is 64.											
<b>level-1</b>	(Optional) Sets the priority for Level 1 independently.											
<b>level-2</b>	(Optional) Sets the priority for Level 2 independently.											
Release	Modification											
4.0(1)	This command was introduced.											

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record														
<p><b>isis priority</b></p> <p>To configure the priority of designated routers, use the <b>isis priority</b> command in interface configuration mode. To reset the default priority, use the <b>no</b> form of this command.</p> <pre>isis priority number-value [level-1   level-2] no isis priority [level-1   level-2]</pre> <p><b>Syntax Description</b></p> <table border="1"> <tr> <td><b>number-value</b></td> <td>Priority of a router and is a number from 0 to 127. The default value is 64.</td> </tr> <tr> <td><b>level-1</b></td> <td>(Optional) Sets the priority for Level 1 independently.</td> </tr> <tr> <td><b>level-2</b></td> <td>(Optional) Sets the priority for Level 2 independently.</td> </tr> </table> <p><b>Defaults</b></p> <p>Priority of 64 Level 1 and Level 2</p> <p><b>Command Modes</b></p> <p>Interface configuration</p> <p><b>Supported User Roles</b></p> <p>network-admin vdc-admin</p> <p><b>Command History</b></p> <table border="1"> <tr> <th>Release</th> <th>Modification</th> </tr> <tr> <td>4.0(1)</td> <td>This command was introduced.</td> </tr> </table> <p><b>Usage Guidelines</b></p> <p>Priorities can be configured for Level 1 and Level 2 independently. Specifying the <b>level-1</b> or <b>level-2</b> keyword resets priority only for Level 1 or Level 2 routing, respectively.</p> <p>The priority is used to determine which router on a LAN will be the designated router or Designated Intermediate System (DIS). The priorities are advertised in the hello packets. The router with the highest priority will become the DIS.</p> <p>In Intermediate System-to-Intermediate System (IS-IS), there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a router with a higher priority comes on line, it will take over the role from the current DIS. In the case of equal priorities, the highest MAC address breaks the tie.</p> <p>This command requires the Enterprise Services license.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 5.x (2010), at L3-397.</p>	<b>number-value</b>	Priority of a router and is a number from 0 to 127. The default value is 64.	<b>level-1</b>	(Optional) Sets the priority for Level 1 independently.	<b>level-2</b>	(Optional) Sets the priority for Level 2 independently.	Release	Modification	4.0(1)	This command was introduced.	<p><b>isis priority</b></p> <p>The <b>isis priority</b> command configures IS-IS router priority for the configuration mode interface. The priority is used to determine which device will be the Designated Intermediate System (DIS). The device with the highest priority will become the DIS.</p> <p>In IS-IS, there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a device with a higher priority comes on line, it will take over the role from the current DIS.</p> <p>The <b>no isis priority</b> and <b>default isis priority</b> commands restore the default priority (64) on the configuration mode interface.</p> <table> <tr> <td><b>Platform</b></td> <td>all</td> </tr> <tr> <td><b>Command Mode</b></td> <td>Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration</td> </tr> </table> <p><b>Command Syntax</b></p> <pre>isis priority priority_level no isis priority default isis priority</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>priority_level</b> priority level. Value ranges from 0 to 127. Default value is 64.</li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1690.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1450.</p>	<b>Platform</b>	all	<b>Command Mode</b>	Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration	Dkt. 419-10 at PDF p. 332
<b>number-value</b>	Priority of a router and is a number from 0 to 127. The default value is 64.															
<b>level-1</b>	(Optional) Sets the priority for Level 1 independently.															
<b>level-2</b>	(Optional) Sets the priority for Level 2 independently.															
Release	Modification															
4.0(1)	This command was introduced.															
<b>Platform</b>	all															
<b>Command Mode</b>	Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration															

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record														
<p><b>isis priority</b></p> <p>To configure the priority of designated routers, use the <b>isis priority</b> command in interface configuration mode. To reset the default priority, use the <b>no</b> form of this command.</p> <pre>isis priority number-value [level-1   level-2] no isis priority [level-1   level-2]</pre> <p><b>Syntax Description</b></p> <table border="1"> <tr> <td><b>number-value</b></td> <td>Priority of a router and is a number from 0 to 127. The default value is 64.</td> </tr> <tr> <td><b>level-1</b></td> <td>(Optional) Sets the priority for Level 1 independently.</td> </tr> <tr> <td><b>level-2</b></td> <td>(Optional) Sets the priority for Level 2 independently.</td> </tr> </table> <p><b>Defaults</b></p> <p>Priority of 64 Level 1 and Level 2</p> <p><b>Command Modes</b></p> <p>Interface configuration</p> <p><b>Supported User Roles</b></p> <p>network-admin vdc-admin</p> <p><b>Command History</b></p> <table border="1"> <tr> <th>Release</th> <th>Modification</th> </tr> <tr> <td>4.0(1)</td> <td>This command was introduced.</td> </tr> </table> <p><b>Usage Guidelines</b></p> <p>Priorities can be configured for Level 1 and Level 2 independently. Specifying the <b>level-1</b> or <b>level-2</b> keyword resets priority only for Level 1 or Level 2 routing, respectively.</p> <p>The priority is used to determine which router on a LAN will be the designated router or Designated Intermediate System (DIS). The priorities are advertised in the hello packets. The router with the highest priority will become the DIS.</p> <p>In Intermediate System-to-Intermediate System (IS-IS), there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a router with a higher priority comes on line, it will take over the role from the current DIS. In the case of equal priorities, the highest MAC address breaks the tie.</p> <p>This command requires the Enterprise Services license.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 4.x (2008), at L3-232.</p>	<b>number-value</b>	Priority of a router and is a number from 0 to 127. The default value is 64.	<b>level-1</b>	(Optional) Sets the priority for Level 1 independently.	<b>level-2</b>	(Optional) Sets the priority for Level 2 independently.	Release	Modification	4.0(1)	This command was introduced.	<p><b>isis priority</b></p> <p>The <b>isis priority</b> command configures IS-IS router priority for the configuration mode interface. The priority is used to determine which device will be the Designated Intermediate System (DIS). The device with the highest priority will become the DIS.</p> <p>In IS-IS, there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a device with a higher priority comes on line, it will take over the role from the current DIS.</p> <p>The <b>no isis priority</b> and <b>default isis priority</b> commands restore the default priority (64) on the configuration mode interface.</p> <table> <tr> <td><b>Platform</b></td> <td>all</td> </tr> <tr> <td><b>Command Mode</b></td> <td>Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration</td> </tr> </table> <p><b>Command Syntax</b></p> <pre>isis priority priority_level no isis priority default isis priority</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>priority_level</b> priority level. Value ranges from 0 to 127. Default value is 64.</li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1690.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1450.</p>	<b>Platform</b>	all	<b>Command Mode</b>	Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration	Dkt. 419-10 at PDF p. 333
<b>number-value</b>	Priority of a router and is a number from 0 to 127. The default value is 64.															
<b>level-1</b>	(Optional) Sets the priority for Level 1 independently.															
<b>level-2</b>	(Optional) Sets the priority for Level 2 independently.															
Release	Modification															
4.0(1)	This command was introduced.															
<b>Platform</b>	all															
<b>Command Mode</b>	Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration															

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record										
<p><b>isis priority</b></p> <p>To configure the priority of designated routers, use the <b>isis priority</b> command in interface configuration mode. To reset the default priority, use the <b>no</b> form of this command.</p> <pre>isis priority number-value [level-1   level-2] no isis priority [level-1   level-2]</pre> <p><b>Syntax Description</b></p> <table border="1"> <tr> <td><b>number-value</b></td> <td>Priority of a router and is a number from 0 to 127. The default value is 64.</td> </tr> <tr> <td><b>level-1</b></td> <td>(Optional) Sets the priority for Level 1 independently.</td> </tr> <tr> <td><b>level-2</b></td> <td>(Optional) Sets the priority for Level 2 independently.</td> </tr> </table> <p><b>Defaults</b></p> <p>Priority of 64 Level 1 and Level 2</p> <p><b>Command Modes</b></p> <p>Interface configuration</p> <p><b>SupportedUserRoles</b></p> <p>network-admin vdc-admin</p> <p><b>Command History</b></p> <table border="1"> <tr> <th>Release</th> <th>Modification</th> </tr> <tr> <td>4.0(1)</td> <td>This command was introduced.</td> </tr> </table> <p><b>Usage Guidelines</b></p> <p>Priorities can be configured for Level 1 and Level 2 independently. Specifying the <b>level-1</b> or <b>level-2</b> keyword resets priority only for Level 1 or Level 2 routing, respectively.</p> <p>The priority is used to determine which router on a LAN will be the designated router or Designated Intermediate System (DIS). The priorities are advertised in the hello packets. The router with the highest priority will become the DIS.</p> <p>In Intermediate System-to-Intermediate System (IS-IS), there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a router with a higher priority comes on line, it will take over the role from the current DIS. In the case of equal priorities, the highest MAC address breaks the tie.</p> <p>This command requires the Enterprise Services license.</p> <p>Cisco IOS IP Routing: ISIS Command Reference (2009), at IRS-63.</p>	<b>number-value</b>	Priority of a router and is a number from 0 to 127. The default value is 64.	<b>level-1</b>	(Optional) Sets the priority for Level 1 independently.	<b>level-2</b>	(Optional) Sets the priority for Level 2 independently.	Release	Modification	4.0(1)	This command was introduced.	<p><b>isis priority</b></p> <p>The <b>isis priority</b> command configures IS-IS router priority for the configuration mode interface.</p> <p>The priority is used to determine which device will be the Designated Intermediate System (DIS). The device with the highest priority will become the DIS.</p> <p>In IS-IS, there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a device with a higher priority comes on line, it will take over the role from the current DIS.</p> <p>The no <b>isis priority</b> and default <b>isis priority</b> commands restore the default priority (64) on the configuration mode interface.</p> <p><b>Platform</b> all  <b>Command Mode</b> Interface-Ethernet Configuration Interface-Loopback Configuration Interface-Port-channel Configuration Interface-VLAN Configuration</p> <p><b>Command Syntax</b></p> <pre>isis priority priority_level no isis priority default isis priority</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li><b>priority_level</b> priority level. Value ranges from 0 to 127. Default value is 64.</li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1690.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1450.</p>	Dkt. 419-10 at PDF p. 334
<b>number-value</b>	Priority of a router and is a number from 0 to 127. The default value is 64.											
<b>level-1</b>	(Optional) Sets the priority for Level 1 independently.											
<b>level-2</b>	(Optional) Sets the priority for Level 2 independently.											
Release	Modification											
4.0(1)	This command was introduced.											

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record										
<p><b>log-adjacency-changes (IS-IS)</b></p> <p>To enable the router to send a syslog message when an Intermediate System-to-Intermediate System Intradomain Routing Protocol (IS-IS) neighbor goes up or down, use the <code>log-adjacency-changes</code> configuration mode command. To disable this function, use the <code>no</code> form of this command.</p> <div style="border: 1px solid red; padding: 2px; display: inline-block;"> <code>log-adjacency-changes</code> </div> <div style="border: 1px solid red; padding: 2px; display: inline-block;"> <code>no log-adjacency-changes</code> </div> <p><b>Syntax Description</b> This command has no arguments or keywords.</p> <p><b>Defaults</b> This command is enabled by default.</p> <p><b>Command Modes</b> Router configuration VRF configuration</p> <p><b>SupportedUserRoles</b> network-admin vdc-admin</p> <p><b>Command History</b></p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Release</th> <th style="text-align: left;">Modification</th> </tr> </thead> <tbody> <tr> <td>4.0(1)</td> <td>This command was introduced.</td> </tr> </tbody> </table> <p><b>Usage Guidelines</b> The <code>log-adjacency-changes</code> command is on by default but only up/down (full/down) events are reported.</p> <p><b>Examples</b> This example configures the router to send a syslog message when an IS-IS neighbor state changes:</p> <pre>switch(config)# router isis Osiris switch(config-router)# log-adjacency-changes</pre> <p><b>Related Commands</b></p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Command</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td><code>feature isis</code></td> <td>Enables IS-IS on the router.</td> </tr> <tr> <td><code>router isis</code></td> <td>Enables IS-IS.</td> </tr> </tbody> </table> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 138.</p>	Release	Modification	4.0(1)	This command was introduced.	Command	Description	<code>feature isis</code>	Enables IS-IS on the router.	<code>router isis</code>	Enables IS-IS.	<p><b>log-adjacency-changes (IS-IS)</b></p> <p>The <code>log-adjacency-changes</code> command configures the switch to send syslog messages either when it detects IS-IS link state changes or when it detects that a neighbor has gone up or down. Log message sending is disabled by default.</p> <p>The default option is active when <code>running-config</code> does not contain any form of the command. Entering the command in any form replaces the previous command state in <code>running-config</code>.</p> <p>Platform all Command Mode Router-IS-IS Configuration</p> <p><b>Command Syntax</b></p> <div style="border: 1px solid red; padding: 2px; display: inline-block;"> <code>log-adjacency-changes</code> </div> <div style="border: 1px solid red; padding: 2px; display: inline-block;"> <code>no log-adjacency-changes</code> </div> <div style="border: 1px solid red; padding: 2px; display: inline-block;"> <code>default log-adjacency-changes</code> </div> <p><b>Examples</b></p> <ul style="list-style-type: none"> <li>These commands configure the switch to send a syslog message when a neighbor goes up or down.</li> </ul> <pre>switch(config)#router isis Osiris switch(config-router-isis)#log-adjacency-changes switch(config-router-isis)# </pre> <ul style="list-style-type: none"> <li>These commands configure not to log the peer changes.</li> </ul> <pre>switch(config)#router isis Osiris switch(config-router-isis)#no log-adjacency-changes switch(config-router-isis)# </pre> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1692.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1452.</p>	Dkt. 419-10 at PDF p. 335
Release	Modification											
4.0(1)	This command was introduced.											
Command	Description											
<code>feature isis</code>	Enables IS-IS on the router.											
<code>router isis</code>	Enables IS-IS.											

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record										
<p><b>log-adjacency-changes (IS-IS)</b></p> <p>To enable the router to send a syslog message when an Intermediate System-to-Intermediate System Intradomain Routing Protocol (IS-IS) neighbor goes up or down, use the <code>log-adjacency-changes</code> configuration mode command. To disable this function, use the <code>no</code> form of this command.</p> <div style="border: 1px solid red; padding: 2px; display: inline-block;"> <code>log-adjacency-changes</code> </div> <div style="border: 1px solid red; padding: 2px; display: inline-block;"> <code>no log-adjacency-changes</code> </div> <p><b>Syntax Description</b> This command has no arguments or keywords.</p> <p><b>Defaults</b> This command is enabled by default.</p> <p><b>Command Modes</b> Router configuration VRF configuration</p> <p><b>SupportedUserRoles</b> network-admin vdc-admin</p> <p><b>Command History</b></p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Release</th> <th style="text-align: left;">Modification</th> </tr> </thead> <tbody> <tr> <td style="text-align: left;">4.0(1)</td> <td>This command was introduced.</td> </tr> </tbody> </table> <p><b>Usage Guidelines</b> The <code>log-adjacency-changes</code> command is on by default but only up/down (full/down) events are reported.</p> <p><b>Examples</b> This example configures the router to send a syslog message when an IS-IS neighbor state changes:</p> <pre>switch(config)# router isis Osiris switch(config-router)# log-adjacency-changes</pre> <p><b>Related Commands</b></p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Command</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td style="text-align: left;"><code>feature isis</code></td> <td>Enables IS-IS on the router.</td> </tr> <tr> <td style="text-align: left;"><code>router isis</code></td> <td>Enables IS-IS.</td> </tr> </tbody> </table> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 5.x (2010), at L3-403.</p>	Release	Modification	4.0(1)	This command was introduced.	Command	Description	<code>feature isis</code>	Enables IS-IS on the router.	<code>router isis</code>	Enables IS-IS.	<p><b>log-adjacency-changes (IS-IS)</b></p> <p>The <code>log-adjacency-changes</code> command configures the switch to send syslog messages either when it detects IS-IS link state changes or when it detects that a neighbor has gone up or down. Log message sending is disabled by default.</p> <p>The default option is active when <code>running-config</code> does not contain any form of the command. Entering the command in any form replaces the previous command state in <code>running-config</code>.</p> <p>Platform all Command Mode Router-IS-IS Configuration</p> <p><b>Command Syntax</b></p> <div style="border: 1px solid red; padding: 2px; display: inline-block;"> <code>log-adjacency-changes</code> </div> <div style="border: 1px solid red; padding: 2px; display: inline-block;"> <code>no log-adjacency-changes</code> </div> <div style="border: 1px solid red; padding: 2px; display: inline-block;"> <code>default log-adjacency-changes</code> </div> <p><b>Examples</b></p> <ul style="list-style-type: none"> <li>These commands configure the switch to send a syslog message when a neighbor goes up or down.</li> </ul> <pre>switch(config)#router isis Osiris switch(config-router-isis)#log-adjacency-changes switch(config-router-isis)# </pre> <ul style="list-style-type: none"> <li>These commands configure not to log the peer changes.</li> </ul> <pre>switch(config)#router isis Osiris switch(config-router-isis)#no log-adjacency-changes switch(config-router-isis)# </pre> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1692.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1452.</p>	Dkt. 419-10 at PDF p. 336
Release	Modification											
4.0(1)	This command was introduced.											
Command	Description											
<code>feature isis</code>	Enables IS-IS on the router.											
<code>router isis</code>	Enables IS-IS.											

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record										
<p><b>log-adjacency-changes (IS-IS)</b></p> <p>To enable the router to send a syslog message when an Intermediate System-to-Intermediate System Intradomain Routing Protocol (IS-IS) neighbor goes up or down, use the <code>log-adjacency-changes</code> configuration mode command. To disable this function, use the <code>no</code> form of this command.</p> <div style="border: 1px solid red; padding: 2px; display: inline-block;"> <code>log-adjacency-changes</code> </div> <div style="border: 1px solid red; padding: 2px; display: inline-block;"> <code>no log-adjacency-changes</code> </div> <p><b>Syntax Description</b> This command has no arguments or keywords.</p> <p><b>Defaults</b> This command is enabled by default.</p> <p><b>Command Modes</b> Router configuration VRF configuration</p> <p><b>SupportedUserRoles</b> network-admin vdc-admin</p> <p><b>Command History</b></p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Release</th> <th style="text-align: left;">Modification</th> </tr> </thead> <tbody> <tr> <td style="text-align: left;">4.0(1)</td> <td>This command was introduced.</td> </tr> </tbody> </table> <p><b>Usage Guidelines</b> The <code>log-adjacency-changes</code> command is on by default but only up/down (full/down) events are reported.</p> <p><b>Examples</b> This example configures the router to send a syslog message when an IS-IS neighbor state changes:</p> <pre>switch(config)# router isis Osiris switch(config-router)# log-adjacency-changes</pre> <p><b>Related Commands</b></p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Command</th> <th style="text-align: left;">Description</th> </tr> </thead> <tbody> <tr> <td style="text-align: left;"><code>feature isis</code></td> <td>Enables IS-IS on the router.</td> </tr> <tr> <td style="text-align: left;"><code>router isis</code></td> <td>Enables IS-IS.</td> </tr> </tbody> </table> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 4.x (2008), at L3-235.</p>	Release	Modification	4.0(1)	This command was introduced.	Command	Description	<code>feature isis</code>	Enables IS-IS on the router.	<code>router isis</code>	Enables IS-IS.	<p><b>log-adjacency-changes (IS-IS)</b></p> <p>The <code>log-adjacency-changes</code> command configures the switch to send syslog messages either when it detects IS-IS link state changes or when it detects that a neighbor has gone up or down. Log message sending is disabled by default.</p> <p>The default option is active when <code>running-config</code> does not contain any form of the command. Entering the command in any form replaces the previous command state in <code>running-config</code>.</p> <p>Platform all Command Mode Router-IS-IS Configuration</p> <p><b>Command Syntax</b></p> <div style="border: 1px solid red; padding: 2px; display: inline-block;"> <code>log-adjacency-changes</code> </div> <div style="border: 1px solid red; padding: 2px; display: inline-block;"> <code>no log-adjacency-changes</code> </div> <div style="border: 1px solid red; padding: 2px; display: inline-block;"> <code>default log-adjacency-changes</code> </div> <p><b>Examples</b></p> <ul style="list-style-type: none"> <li>These commands configure the switch to send a syslog message when a neighbor goes up or down.</li> </ul> <pre>switch(config)#router isis Osiris switch(config-router-isis)#log-adjacency-changes switch(config-router-isis)# </pre> <ul style="list-style-type: none"> <li>These commands configure not to log the peer changes.</li> </ul> <pre>switch(config)#router isis Osiris switch(config-router-isis)#no log-adjacency-changes switch(config-router-isis)# </pre> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1692.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1452.</p>	Dkt. 419-10 at PDF p. 337
Release	Modification											
4.0(1)	This command was introduced.											
Command	Description											
<code>feature isis</code>	Enables IS-IS on the router.											
<code>router isis</code>	Enables IS-IS.											

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record																				
<p><b>max-metric router-lsa (OSPF)</b></p> <p>To configure the Open Shortest Path First (OSPF) protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their shortest path first (SPF) calculations, use the <b>max-metric router-lsa</b> command. To disable the advertisement of a maximum metric, use the <b>no</b> form of this command.</p> <pre>max-metric router-lsa [external] lsa [max-metric-value] [include-stub] [on-startup] [seconds] [wait-for bgp tag] [summary-lsa [max-metric-value]] no max-metric router-lsa [external] lsa [max-metric-value] [include-stub] [on-startup] [seconds] [wait-for bgp tag] [summary-lsa [max-metric-value]]</pre> <table border="1"> <tr> <td><b>external-lsa</b></td> <td>Specifies the external LSA's.</td> </tr> <tr> <td><b>max-metric-value</b></td> <td>(Optional) Specifies the max-metric values for external LSA's. The range is 1-65535.</td> </tr> <tr> <td><b>include-stub</b></td> <td>Advertises the max-metric for stub links.</td> </tr> <tr> <td><b>on-startup</b></td> <td>(Optional) Configures the router to advertise a maximum metric at startup.</td> </tr> <tr> <td><b>seconds</b></td> <td>(Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds.</td> </tr> <tr> <td><b>wait-for bgp tag</b></td> <td>(Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.</td> </tr> <tr> <td><b>summary-lsa</b></td> <td>Specifies the summary LSA's.</td> </tr> <tr> <td><b>max-metric-value</b></td> <td>(Optional) Specifies the max-metric value for summary LSAs. The range is from 1-65535.</td> </tr> </table> <p><b>Defaults</b></p> <p>Originates router link-state advertisements (LSAs) with normal link metrics.</p> <p><b>Command Modes</b></p> <p>Router configuration Router VRF configuration</p> <p><b>SupportedUserRoles</b></p> <p>network-admin vdc-admin</p> <p><b>Command History</b></p> <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>4.0(1)</td> <td>This command was introduced.</td> </tr> </tbody> </table> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 194.</p>	<b>external-lsa</b>	Specifies the external LSA's.	<b>max-metric-value</b>	(Optional) Specifies the max-metric values for external LSA's. The range is 1-65535.	<b>include-stub</b>	Advertises the max-metric for stub links.	<b>on-startup</b>	(Optional) Configures the router to advertise a maximum metric at startup.	<b>seconds</b>	(Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds.	<b>wait-for bgp tag</b>	(Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.	<b>summary-lsa</b>	Specifies the summary LSA's.	<b>max-metric-value</b>	(Optional) Specifies the max-metric value for summary LSAs. The range is from 1-65535.	Release	Modification	4.0(1)	This command was introduced.	<p><b>max-metric router-lsa (OSPFv2)</b></p> <p>The <b>max-metric router-lsa</b> command allows the OSPF protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their SPF calculations.</p> <p>The <b>no max-metric router-lsa</b> and <b>default max-metric router-lsa</b> commands disable the advertisement of a maximum metric.</p> <p>Platform all Command Mode Router-OSPF Configuration</p> <p><b>Command Syntax</b></p> <pre>max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY] no max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY] default max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY]</pre> <p>All parameters can be placed in any order.</p> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>EXTERNAL</b> advertised metric value. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; Metric is set to the default value of 1.</li> <li>— <b>external-lsa</b> Configures the router to override the External LSA / NSSA-External metric with the maximum metric value.</li> <li>— <b>external-lsa &lt;1 to 16777215&gt;</b> The configurable range is from 1 to 0xFFFFFFF. The default value is 0xFF0000. This range can be used with external LSA, summary LSA extensions to indicate the respective metric you want with the LSA.</li> </ul> </li> <li>• <b>STUB</b> advertised metric type. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; Metric type is set to the default value of 2.</li> <li>— <b>include-stub</b> Advertises stub links in router-LSA with the max-metric value (0xFFFF).</li> </ul> </li> <li>• <b>STARTUP</b> limit scope of LSAs. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; LSA can be translated</li> <li>— <b>on-startup</b> Configures the router to advertise a maximum metric at startup (only valid in no and default command formats).</li> <li>— <b>on-startup wait-for-bgp</b> Configures the router to advertise a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.</li> <li>— <b>on-startup &lt;5 to 86400&gt;</b> Sets the maximum metric temporarily after a reboot to originate router-LSAs with the max-metric value. <b>wait-for-bgp</b> or an on-start time value is not included in no and default commands.</li> </ul> </li> <li>• <b>SUMMARY</b> advertised metric value. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; Metric is set to the default value of 1.</li> <li>— <b>summary-lsa</b> Configures the router to override the summary LSA metric with the maximum metric value for both type 3 and type 4 Summary LSAs.</li> <li>— <b>summary-lsa &lt;1 to 16777215&gt;</b> Metric is set to the specified value.</li> </ul> </li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1439.</p>	Dkt. 419-10 at PDF p. 338
<b>external-lsa</b>	Specifies the external LSA's.																					
<b>max-metric-value</b>	(Optional) Specifies the max-metric values for external LSA's. The range is 1-65535.																					
<b>include-stub</b>	Advertises the max-metric for stub links.																					
<b>on-startup</b>	(Optional) Configures the router to advertise a maximum metric at startup.																					
<b>seconds</b>	(Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds.																					
<b>wait-for bgp tag</b>	(Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.																					
<b>summary-lsa</b>	Specifies the summary LSA's.																					
<b>max-metric-value</b>	(Optional) Specifies the max-metric value for summary LSAs. The range is from 1-65535.																					
Release	Modification																					
4.0(1)	This command was introduced.																					

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record																				
<p><b>max-metric router-lsa (OSPF)</b></p> <p>To configure the Open Shortest Path First (OSPF) protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their shortest path first (SPF) calculations, use the <b>max-metric router-lsa</b> command. To disable the advertisement of a maximum metric, use the <b>no</b> form of this command.</p> <pre>max-metric router-lsa [external-lsa [max-metric-value]] [include-stub] [on-startup] [seconds] [wait-for bgp tag]] [summary-lsa [max-metric-value]] no max-metric router-lsa [external-lsa [max-metric-value]] [include-stub] [on-startup] [seconds] [wait-for bgp tag]] [summary-lsa [max-metric-value]]</pre> <table border="1"> <tr> <td><b>external-lsa</b></td> <td>Specifies the external LSA's.</td> </tr> <tr> <td><b>max-metric-value</b></td> <td>(Optional) Specifies the max-metric values for external LSA's. The range is 1-65535.</td> </tr> <tr> <td><b>include-stub</b></td> <td>Advertises the max-metric for stub links.</td> </tr> <tr> <td><b>on-startup</b></td> <td>(Optional) Configures the router to advertise a maximum metric at startup.</td> </tr> <tr> <td><b>seconds</b></td> <td>(Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds.</td> </tr> <tr> <td><b>wait-for bgp tag</b></td> <td>(Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.</td> </tr> <tr> <td><b>summary-lsa</b></td> <td>Specifies the summary LSA's.</td> </tr> <tr> <td><b>max-metric-value</b></td> <td>(Optional) Specifies the max-metric value for summary LSAs. The range is from 1-65535.</td> </tr> </table> <p><b>Defaults</b></p> <p>Originates router link-state advertisements (LSAs) with normal link metrics.</p> <p><b>Command Modes</b></p> <p>Router configuration Router VRF configuration</p> <p><b>SupportedUserRoles</b></p> <p>network-admin vdc-admin</p> <p><b>Command History</b></p> <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>4.0(1)</td> <td>This command was introduced.</td> </tr> </tbody> </table>	<b>external-lsa</b>	Specifies the external LSA's.	<b>max-metric-value</b>	(Optional) Specifies the max-metric values for external LSA's. The range is 1-65535.	<b>include-stub</b>	Advertises the max-metric for stub links.	<b>on-startup</b>	(Optional) Configures the router to advertise a maximum metric at startup.	<b>seconds</b>	(Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds.	<b>wait-for bgp tag</b>	(Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.	<b>summary-lsa</b>	Specifies the summary LSA's.	<b>max-metric-value</b>	(Optional) Specifies the max-metric value for summary LSAs. The range is from 1-65535.	Release	Modification	4.0(1)	This command was introduced.	<p><b>max-metric router-lsa (OSPFv2)</b></p> <p>The <b>max-metric router-lsa</b> command allows the OSPF protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their SPF calculations.</p> <p>The <b>no max-metric router-lsa</b> and <b>default max-metric router-lsa</b> commands disable the advertisement of a maximum metric.</p> <p>Platform all Command Mode Router-OSPF Configuration</p> <p><b>Command Syntax</b></p> <pre>max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY] no max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY] default max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY]</pre> <p>All parameters can be placed in any order.</p> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>EXTERNAL</b> advertised metric value. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; Metric is set to the default value of 1.</li> <li>— <b>external-lsa</b> Configures the router to override the External LSA / NSSA-External metric with the maximum metric value.</li> <li>— <b>external-lsa &lt;1 to 16777215&gt;</b> The configurable range is from 1 to 0xFFFFFFF. The default value is 0xFF0000. This range can be used with external LSA, summary LSA extensions to indicate the respective metric you want with the LSA.</li> </ul> </li> <li>• <b>STUB</b> advertised metric type. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; Metric type is set to the default value of 2.</li> <li>— <b>include-stub</b> Advertises stub links in router-LSA with the max-metric value (0xFFFF).</li> </ul> </li> <li>• <b>STARTUP</b> limit scope of LSAs. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; LSA can be translated</li> <li>— <b>on-startup</b> Configures the router to advertise a maximum metric at startup (only valid in no and default command formats).</li> <li>— <b>on-startup wait-for-bgp</b> Configures the router to advertise a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.</li> <li>— <b>on-startup &lt;5 to 86400&gt;</b> Sets the maximum metric temporarily after a reboot to originate router-LSAs with the max-metric value. <b>wait-for-bgp</b> or an on-start time value is not included in no and default commands.</li> </ul> </li> <li>• <b>SUMMARY</b> advertised metric value. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; Metric is set to the default value of 1.</li> <li>— <b>summary-lsa</b> Configures the router to override the summary LSA metric with the maximum metric value for both type 3 and type 4 Summary LSAs.</li> <li>— <b>summary-lsa &lt;1 to 16777215&gt;</b> Metric is set to the specified value.</li> </ul> </li> </ul>	Dkt. 419-10 at PDF p. 339
<b>external-lsa</b>	Specifies the external LSA's.																					
<b>max-metric-value</b>	(Optional) Specifies the max-metric values for external LSA's. The range is 1-65535.																					
<b>include-stub</b>	Advertises the max-metric for stub links.																					
<b>on-startup</b>	(Optional) Configures the router to advertise a maximum metric at startup.																					
<b>seconds</b>	(Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds.																					
<b>wait-for bgp tag</b>	(Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.																					
<b>summary-lsa</b>	Specifies the summary LSA's.																					
<b>max-metric-value</b>	(Optional) Specifies the max-metric value for summary LSAs. The range is from 1-65535.																					
Release	Modification																					
4.0(1)	This command was introduced.																					

Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 5.x (2010), at L3-457

Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1439.

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record																				
<p><b>max-metric router-lsa (OSPF)</b></p> <p>To configure the Open Shortest Path First (OSPF) protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their shortest path first (SPF) calculations, use the <b>max-metric router-lsa</b> command. To disable the advertisement of a maximum metric, use the <b>no</b> form of this command.</p> <pre>max-metric router-lsa [external] lsa [max-metric-value] [include-stub] [on-startup] [seconds] [wait-for bgp tag] [summary-lsa [max-metric-value]] no max-metric router-lsa [external] lsa [max-metric-value] [include-stub] [on-startup] [seconds] [wait-for bgp tag] [summary-lsa [max-metric-value]]</pre> <table border="1"> <tr> <td><b>external-lsa</b></td> <td>Specifies the external LSA's.</td> </tr> <tr> <td><b>max-metric-value</b></td> <td>(Optional) Specifies the max-metric values for external LSA's. The range is 1-65535.</td> </tr> <tr> <td><b>include-stub</b></td> <td>Advertises the max-metric for stub links.</td> </tr> <tr> <td><b>on-startup</b></td> <td>(Optional) Configures the router to advertise a maximum metric at startup.</td> </tr> <tr> <td><b>seconds</b></td> <td>(Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds.</td> </tr> <tr> <td><b>wait-for bgp tag</b></td> <td>(Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.</td> </tr> <tr> <td><b>summary-lsa</b></td> <td>Specifies the summary LSA's.</td> </tr> <tr> <td><b>max-metric-value</b></td> <td>(Optional) Specifies the max-metric value for summary LSAs. The range is from 1-65535.</td> </tr> </table> <p><b>Defaults</b></p> <p>Originates router link-state advertisements (LSAs) with normal link metrics.</p> <p><b>Command Modes</b></p> <p>Router configuration Router VRF configuration</p> <p><b>SupportedUserRoles</b></p> <p>network-admin vdc-admin</p> <p><b>Command History</b></p> <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>4.0(1)</td> <td>This command was introduced.</td> </tr> </tbody> </table> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 4.x (2008), at L3-272</p>	<b>external-lsa</b>	Specifies the external LSA's.	<b>max-metric-value</b>	(Optional) Specifies the max-metric values for external LSA's. The range is 1-65535.	<b>include-stub</b>	Advertises the max-metric for stub links.	<b>on-startup</b>	(Optional) Configures the router to advertise a maximum metric at startup.	<b>seconds</b>	(Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds.	<b>wait-for bgp tag</b>	(Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.	<b>summary-lsa</b>	Specifies the summary LSA's.	<b>max-metric-value</b>	(Optional) Specifies the max-metric value for summary LSAs. The range is from 1-65535.	Release	Modification	4.0(1)	This command was introduced.	<p><b>max-metric router-lsa (OSPFv2)</b></p> <p>The <b>max-metric router-lsa</b> command allows the OSPF protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their SPF calculations.</p> <p>The no <b>max-metric router-lsa</b> and default <b>max-metric router-lsa</b> commands disable the advertisement of a maximum metric.</p> <p>Platform all Command Mode Router-OSPF Configuration</p> <p><b>Command Syntax</b></p> <pre>max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY] no max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY] default max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY]</pre> <p>All parameters can be placed in any order.</p> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>EXTERNAL</b> advertised metric value. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; Metric is set to the default value of 1.</li> <li>— <b>external-lsa</b> Configures the router to override the External LSA / NSSA-External metric with the maximum metric value.</li> <li>— <b>external-lsa &lt;1 to 16777215&gt;</b> The configurable range is from 1 to 0xFFFFFFF. The default value is 0xFF0000. This range can be used with external LSA, summary LSA extensions to indicate the respective metric you want with the LSA.</li> </ul> </li> <li>• <b>STUB</b> advertised metric type. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; Metric type is set to the default value of 2.</li> <li>— <b>include-stub</b> Advertises stub links in router-LSA with the max-metric value (0xFFFF).</li> </ul> </li> <li>• <b>STARTUP</b> limit scope of LSAs. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; LSA can be translated</li> <li>— <b>on-startup</b> Configures the router to advertise a maximum metric at startup (only valid in no and default command formats).</li> <li>— <b>on-startup wait-for-bgp</b> Configures the router to advertise a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.</li> <li>— <b>on-startup &lt;5 to 86400&gt;</b> Sets the maximum metric temporarily after a reboot to originate router-LSAs with the max-metric value. <b>wait-for-bgp</b> or an on-start time value is not included in no and default commands.</li> </ul> </li> <li>• <b>SUMMARY</b> advertised metric value. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; Metric is set to the default value of 1.</li> <li>— <b>summary-lsa</b> Configures the router to override the summary LSA metric with the maximum metric value for both type 3 and type 4 Summary LSAs.</li> <li>— <b>summary-lsa &lt;1 to 16777215&gt;</b> Metric is set to the specified value.</li> </ul> </li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1439.</p>	Dkt. 419-10 at PDF p. 340
<b>external-lsa</b>	Specifies the external LSA's.																					
<b>max-metric-value</b>	(Optional) Specifies the max-metric values for external LSA's. The range is 1-65535.																					
<b>include-stub</b>	Advertises the max-metric for stub links.																					
<b>on-startup</b>	(Optional) Configures the router to advertise a maximum metric at startup.																					
<b>seconds</b>	(Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds.																					
<b>wait-for bgp tag</b>	(Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.																					
<b>summary-lsa</b>	Specifies the summary LSA's.																					
<b>max-metric-value</b>	(Optional) Specifies the max-metric value for summary LSAs. The range is from 1-65535.																					
Release	Modification																					
4.0(1)	This command was introduced.																					

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record																				
<p><b>max-metric router-lsa (OSPF)</b></p> <p>To configure the Open Shortest Path First (OSPF) protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their shortest path first (SPF) calculations, use the <b>max-metric router-lsa</b> command. To disable the advertisement of a maximum metric, use the <b>no</b> form of this command.</p> <pre>max-metric router-lsa [external-lsa [max-metric-value]] [include-stub] [on-startup] [seconds] [wait-for bgp tag]] [summary-lsa [max-metric-value]] no max-metric router-lsa [external-lsa [max-metric-value]] [include-stub] [on-startup] [seconds] [wait-for bgp tag]] [summary-lsa [max-metric-value]]</pre> <table border="1"> <tr> <td><b>external-lsa</b></td> <td>Specifies the external LSA's.</td> </tr> <tr> <td><b>max-metric-value</b></td> <td>(Optional) Specifies the max-metric values for external LSA's. The range is 1-65535.</td> </tr> <tr> <td><b>include-stub</b></td> <td>Advertises the max-metric for stub links.</td> </tr> <tr> <td><b>on-startup</b></td> <td>(Optional) Configures the router to advertise a maximum metric at startup.</td> </tr> <tr> <td><b>seconds</b></td> <td>(Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds.</td> </tr> <tr> <td><b>wait-for bgp tag</b></td> <td>(Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.</td> </tr> <tr> <td><b>summary-lsa</b></td> <td>Specifies the summary LSA's.</td> </tr> <tr> <td><b>max-metric-value</b></td> <td>(Optional) Specifies the max-metric value for summary LSAs. The range is from 1-65535.</td> </tr> </table> <p><b>Defaults</b></p> <p>Originates router link-state advertisements (LSAs) with normal link metrics.</p> <p><b>Command Modes</b></p> <p>Router configuration Router VRF configuration</p> <p><b>SupportedUserRoles</b></p> <p>network-admin vdc-admin</p> <p><b>Command History</b></p> <table border="1"> <thead> <tr> <th>Release</th> <th>Modification</th> </tr> </thead> <tbody> <tr> <td>4.0(1)</td> <td>This command was introduced.</td> </tr> </tbody> </table>	<b>external-lsa</b>	Specifies the external LSA's.	<b>max-metric-value</b>	(Optional) Specifies the max-metric values for external LSA's. The range is 1-65535.	<b>include-stub</b>	Advertises the max-metric for stub links.	<b>on-startup</b>	(Optional) Configures the router to advertise a maximum metric at startup.	<b>seconds</b>	(Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds.	<b>wait-for bgp tag</b>	(Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.	<b>summary-lsa</b>	Specifies the summary LSA's.	<b>max-metric-value</b>	(Optional) Specifies the max-metric value for summary LSAs. The range is from 1-65535.	Release	Modification	4.0(1)	This command was introduced.	<p><b>max-metric router-lsa (OSPFv2)</b></p> <p>The <b>max-metric router-lsa</b> command allows the OSPF protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their SPF calculations.</p> <p>The <b>no max-metric router-lsa</b> and <b>default max-metric router-lsa</b> commands disable the advertisement of a maximum metric.</p> <p>Platform all Command Mode Router-OSPF Configuration</p> <p><b>Command Syntax</b></p> <pre>max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY] no max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY] default max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY]</pre> <p>All parameters can be placed in any order.</p> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>EXTERNAL</b> advertised metric value. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; Metric is set to the default value of 1.</li> <li>— <b>external-lsa</b> Configures the router to override the External LSA / NSSA-External metric with the maximum metric value.</li> <li>— <b>external-lsa &lt;1 to 16777215&gt;</b> The configurable range is from 1 to 0xFFFFFFF. The default value is 0xFF0000. This range can be used with external LSA, summary LSA extensions to indicate the respective metric you want with the LSA.</li> </ul> </li> <li>• <b>STUB</b> advertised metric type. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; Metric type is set to the default value of 2.</li> <li>— <b>include-stub</b> Advertises stub links in router-LSA with the max-metric value (0xFFFF).</li> </ul> </li> <li>• <b>STARTUP</b> limit scope of LSAs. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; LSA can be translated</li> <li>— <b>on-startup</b> Configures the router to advertise a maximum metric at startup (only valid in no and default command formats).</li> <li>— <b>on-startup wait-for-bgp</b> Configures the router to advertise a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.</li> <li>— <b>on-startup &lt;5 to 86400&gt;</b> Sets the maximum metric temporarily after a reboot to originate router-LSAs with the max-metric value. <b>wait-for-bgp</b> or an on-start time value is not included in no and default commands.</li> </ul> </li> <li>• <b>SUMMARY</b> advertised metric value. Values include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; Metric is set to the default value of 1.</li> <li>— <b>summary-lsa</b> Configures the router to override the summary LSA metric with the maximum metric value for both type 3 and type 4 Summary LSAs.</li> <li>— <b>summary-lsa &lt;1 to 16777215&gt;</b> Metric is set to the specified value.</li> </ul> </li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1439.</p>	Dkt. 419-10 at PDF p. 341
<b>external-lsa</b>	Specifies the external LSA's.																					
<b>max-metric-value</b>	(Optional) Specifies the max-metric values for external LSA's. The range is 1-65535.																					
<b>include-stub</b>	Advertises the max-metric for stub links.																					
<b>on-startup</b>	(Optional) Configures the router to advertise a maximum metric at startup.																					
<b>seconds</b>	(Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds.																					
<b>wait-for bgp tag</b>	(Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds.																					
<b>summary-lsa</b>	Specifies the summary LSA's.																					
<b>max-metric-value</b>	(Optional) Specifies the max-metric value for summary LSAs. The range is from 1-65535.																					
Release	Modification																					
4.0(1)	This command was introduced.																					

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record																																																												
<pre>BGP table version is 10, local router ID is 3.3.3.3 Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, &gt;-best Path type: i-internal, e-external, c-confed, l-local, a-aggregate, r-redist Origin codes: i - IGP, e - EGP, ? - incomplete   - multipath</pre> <table border="1"> <thead> <tr> <th>Network</th><th>Next Hop</th><th>Metric</th><th>LocPrf</th><th>Weight</th><th>Path</th></tr> </thead> <tbody> <tr> <td>* i200.0.1.100/32</td><td>201.0.25.1</td><td>100</td><td>100</td><td>6553601</td><td>i</td></tr> <tr> <td>*&gt;e</td><td>201.0.13.1</td><td></td><td>0</td><td>6553601</td><td>i</td></tr> <tr> <td>* i200.0.2.100/32</td><td>201.0.25.1</td><td>100</td><td>100</td><td>6553601</td><td>i</td></tr> <tr> <td>*&gt;e</td><td>201.0.13.1</td><td></td><td>0</td><td>6553601</td><td>i</td></tr> <tr> <td>*&gt;1200.0.3.100/32</td><td>0.0.0.0</td><td>100</td><td>32768</td><td>i</td><td></td></tr> </tbody> </table> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 401.</p>	Network	Next Hop	Metric	LocPrf	Weight	Path	* i200.0.1.100/32	201.0.25.1	100	100	6553601	i	*>e	201.0.13.1		0	6553601	i	* i200.0.2.100/32	201.0.25.1	100	100	6553601	i	*>e	201.0.13.1		0	6553601	i	*>1200.0.3.100/32	0.0.0.0	100	32768	i		<pre>switch&gt;show ip bgp neighbors 10.14.4.4 advertised-routes regexp _64502_ BGP routing table information for VRF default Router identifier 172.24.78.191, local AS number 64498 Route status codes: s - suppressed, * - valid, &gt; - active, E - ECMP head, e - ECMP s - Stale Origin codes: i - IGP, e - EGP, ? - incomplete AS Path Attributes: Or-ID - Originator ID, C-LST - Cluster List, LL Nexthop - Link Local Nexthop</pre> <table border="1"> <thead> <tr> <th>Network</th><th>Next Hop</th><th>Metric</th><th>LocPref</th><th>Weight</th><th>Path</th></tr> </thead> <tbody> <tr> <td>* &gt; 10.99.31.0/24</td><td>10.88.202.1</td><td>333</td><td>100</td><td>-</td><td>(64502 64503) 99 i</td></tr> <tr> <td>* &gt; 10.99.41.0/24</td><td>10.88.202.1</td><td>333</td><td>100</td><td>-</td><td>(64502 64503) 99 i</td></tr> <tr> <td>* &gt; 10.99.99.0/24</td><td>10.88.202.1</td><td>333</td><td>100</td><td>-</td><td>(64502 64504) 99 i</td></tr> </tbody> </table> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1637.</p>	Network	Next Hop	Metric	LocPref	Weight	Path	* > 10.99.31.0/24	10.88.202.1	333	100	-	(64502 64503) 99 i	* > 10.99.41.0/24	10.88.202.1	333	100	-	(64502 64503) 99 i	* > 10.99.99.0/24	10.88.202.1	333	100	-	(64502 64504) 99 i	Dkt. 419-10 at PDF p. 342
Network	Next Hop	Metric	LocPrf	Weight	Path																																																									
* i200.0.1.100/32	201.0.25.1	100	100	6553601	i																																																									
*>e	201.0.13.1		0	6553601	i																																																									
* i200.0.2.100/32	201.0.25.1	100	100	6553601	i																																																									
*>e	201.0.13.1		0	6553601	i																																																									
*>1200.0.3.100/32	0.0.0.0	100	32768	i																																																										
Network	Next Hop	Metric	LocPref	Weight	Path																																																									
* > 10.99.31.0/24	10.88.202.1	333	100	-	(64502 64503) 99 i																																																									
* > 10.99.41.0/24	10.88.202.1	333	100	-	(64502 64503) 99 i																																																									
* > 10.99.99.0/24	10.88.202.1	333	100	-	(64502 64504) 99 i																																																									
<p><b>show ip bgp neighbors</b></p> <p>To display Border Gateway Protocol (BGP) neighbors, use the <code>show ip bgp neighbors</code> command.</p> <pre>show ip bgp neighbors [addr] advertised-routes [flap-statistics   paths   received-routes   routes [advertised   dampened   received]]   prefix [vrf { all   vrf-name }]</pre> <p><b>Syntax Description</b></p> <table border="1"> <tr> <td><b>addr</b></td> <td>IPv4 address. The format is x.x.x.x</td> </tr> <tr> <td><b>advertised-routes</b></td> <td>(Optional) Displays all the routes advertised to this neighbor.</td> </tr> <tr> <td><b>flap-statistics</b></td> <td>(Optional) Displays flap statistics for the routes received from this neighbor.</td> </tr> <tr> <td><b>paths</b></td> <td>(Optional) Displays AS paths learned from this neighbor.</td> </tr> <tr> <td><b>received-routes</b></td> <td>(Optional) Displays all the routes received from this neighbor.</td> </tr> <tr> <td><b>routes</b></td> <td>(Optional) Displays the routes received or advertised to or from this neighbor.</td> </tr> <tr> <td><b>advertised</b></td> <td>(Optional) Displays all the routes advertised for this neighbor.</td> </tr> <tr> <td><b>dampened</b></td> <td>(Optional) Displays all dampened routes received from this neighbor.</td> </tr> <tr> <td><b>received</b></td> <td>(Optional) Displays all the routes received from this neighbor.</td> </tr> <tr> <td><b>prefix</b></td> <td>(Optional) IPv6 prefix. The format is x.x.x.x/length.</td> </tr> <tr> <td><b>vrf vrf-name</b></td> <td>(Optional) Specifies the virtual router context (VRF) name. The name can be any case-sensitive, alphanumeric string up to 63 characters.</td> </tr> <tr> <td><b>all</b></td> <td>(Optional) Specifies all VRF.</td> </tr> </table> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 466.</p>	<b>addr</b>	IPv4 address. The format is x.x.x.x	<b>advertised-routes</b>	(Optional) Displays all the routes advertised to this neighbor.	<b>flap-statistics</b>	(Optional) Displays flap statistics for the routes received from this neighbor.	<b>paths</b>	(Optional) Displays AS paths learned from this neighbor.	<b>received-routes</b>	(Optional) Displays all the routes received from this neighbor.	<b>routes</b>	(Optional) Displays the routes received or advertised to or from this neighbor.	<b>advertised</b>	(Optional) Displays all the routes advertised for this neighbor.	<b>dampened</b>	(Optional) Displays all dampened routes received from this neighbor.	<b>received</b>	(Optional) Displays all the routes received from this neighbor.	<b>prefix</b>	(Optional) IPv6 prefix. The format is x.x.x.x/length.	<b>vrf vrf-name</b>	(Optional) Specifies the virtual router context (VRF) name. The name can be any case-sensitive, alphanumeric string up to 63 characters.	<b>all</b>	(Optional) Specifies all VRF.	<p><b>show ip bgp neighbors</b></p> <p>The <code>show ip bgp neighbors</code> command displays Border Gateway Protocol (BGP) and TCP session data for a specified IPv4 BGP neighbor, or for all IPv4 BGP neighbors if an address is not included.</p> <p><b>Platform</b> all  <b>Command Mode</b> EXEC</p> <p><b>Command Syntax</b></p> <pre>show ip bgp neighbors [NEIGHBOR ADDR] [VRF INSTANCE]</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <b>NEIGHBOR_ADDR</b> location of neighbors. Options include: <ul style="list-style-type: none"> <li>— <code>&lt;no parameter&gt;</code> command displays information for all IPv4 BGP neighbors.</li> <li>— <code>ipv4 addr</code> command displays information for specified neighbor.</li> </ul> </li> <li>• <b>VRF_INSTANCE</b> specifies VRF instances. <ul style="list-style-type: none"> <li>— <code>&lt;no parameter&gt;</code> displays routing table for context-active VRF.</li> <li>— <code>vrf vrf_name</code> displays routing table for the specified VRF.</li> <li>— <code>vrf all</code> displays routing table for all VRFs.</li> <li>— <code>vrf default</code> displays routing table for default VRF.</li> </ul> </li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1632.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1402; Arista User Manual, v. 4.11.1 (1/11/13), at 1148; Arista User Manual v. 4.10.3 (10/22/12), at 959.</p>	Dkt. 419-10 at PDF p. 342																																				
<b>addr</b>	IPv4 address. The format is x.x.x.x																																																													
<b>advertised-routes</b>	(Optional) Displays all the routes advertised to this neighbor.																																																													
<b>flap-statistics</b>	(Optional) Displays flap statistics for the routes received from this neighbor.																																																													
<b>paths</b>	(Optional) Displays AS paths learned from this neighbor.																																																													
<b>received-routes</b>	(Optional) Displays all the routes received from this neighbor.																																																													
<b>routes</b>	(Optional) Displays the routes received or advertised to or from this neighbor.																																																													
<b>advertised</b>	(Optional) Displays all the routes advertised for this neighbor.																																																													
<b>dampened</b>	(Optional) Displays all dampened routes received from this neighbor.																																																													
<b>received</b>	(Optional) Displays all the routes received from this neighbor.																																																													
<b>prefix</b>	(Optional) IPv6 prefix. The format is x.x.x.x/length.																																																													
<b>vrf vrf-name</b>	(Optional) Specifies the virtual router context (VRF) name. The name can be any case-sensitive, alphanumeric string up to 63 characters.																																																													
<b>all</b>	(Optional) Specifies all VRF.																																																													

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record																								
<p><b>show ip bgp neighbors</b></p> <p>To display Border Gateway Protocol (BGP) neighbors, use the <code>show ip bgp neighbors</code> command.</p> <pre>show ip bgp neighbors [addr] [advertised-routes] [flap-statistics] [paths] [received-routes] [routes] [advertised] [damped] [received]] [prefix] [vrf] [ all   vrf-name ]</pre> <p><b>Syntax Description</b></p> <table border="1"> <tr> <td><code>addr</code></td> <td>IPv4 address. The format is x.x.x.x</td> </tr> <tr> <td><code>advertised-routes</code></td> <td>(Optional) Displays all the routes advertised to this neighbor.</td> </tr> <tr> <td><code>flap-statistics</code></td> <td>(Optional) Displays flap statistics for the routes received from this neighbor.</td> </tr> <tr> <td><code>paths</code></td> <td>(Optional) Displays AS paths learned from this neighbor.</td> </tr> <tr> <td><code>received-routes</code></td> <td>(Optional) Displays all the routes received from this neighbor.</td> </tr> <tr> <td><code>routes</code></td> <td>(Optional) Displays the routes received or advertised to or from this neighbor.</td> </tr> <tr> <td><code>advertised</code></td> <td>(Optional) Displays all the routes advertised for this neighbor.</td> </tr> <tr> <td><code>damped</code></td> <td>(Optional) Displays all dampened routes received from this neighbor.</td> </tr> <tr> <td><code>received</code></td> <td>(Optional) Displays all the routes received from this neighbor.</td> </tr> <tr> <td><code>prefix</code></td> <td>(Optional) IPv6 prefix. The format is x.x.x.x/length.</td> </tr> <tr> <td><code>vrf vrf-name</code></td> <td>(Optional) Specifies the virtual router context (VRF) name. The name can be any case-sensitive, alphanumeric string up to 63 characters.</td> </tr> <tr> <td><code>all</code></td> <td>(Optional) Specifies all VRF.</td> </tr> </table> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 5.x (2010), at L3-686.</p>	<code>addr</code>	IPv4 address. The format is x.x.x.x	<code>advertised-routes</code>	(Optional) Displays all the routes advertised to this neighbor.	<code>flap-statistics</code>	(Optional) Displays flap statistics for the routes received from this neighbor.	<code>paths</code>	(Optional) Displays AS paths learned from this neighbor.	<code>received-routes</code>	(Optional) Displays all the routes received from this neighbor.	<code>routes</code>	(Optional) Displays the routes received or advertised to or from this neighbor.	<code>advertised</code>	(Optional) Displays all the routes advertised for this neighbor.	<code>damped</code>	(Optional) Displays all dampened routes received from this neighbor.	<code>received</code>	(Optional) Displays all the routes received from this neighbor.	<code>prefix</code>	(Optional) IPv6 prefix. The format is x.x.x.x/length.	<code>vrf vrf-name</code>	(Optional) Specifies the virtual router context (VRF) name. The name can be any case-sensitive, alphanumeric string up to 63 characters.	<code>all</code>	(Optional) Specifies all VRF.	<p><b>show ip bgp neighbors</b></p> <p>The <code>show ip bgp neighbors</code> command displays Border Gateway Protocol (BGP) and TCP session data for a specified IPv4 BGP neighbor, or for all IPv4 BGP neighbors if an address is not included.</p> <p>Platform all Command Mode EXEC</p> <p><b>Command Syntax</b></p> <pre>show ip bgp neighbors [NEIGHBOR ADDR] [VRF INSTANCE]</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li>• <code>NEIGHBOR_ADDR</code> location of neighbors. Options include: <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; command displays information for all IPv4 BGP neighbors.</li> <li>— <code>ipv4 addr</code> command displays information for specified neighbor.</li> </ul> </li> <li>• <code>VRF_INSTANCE</code> specifies VRF instances. <ul style="list-style-type: none"> <li>— &lt;no parameter&gt; displays routing table for context-active VRF.</li> <li>— <code>vrf vrf_name</code> displays routing table for the specified VRF.</li> <li>— <code>vrf all</code> displays routing table for all VRFs.</li> <li>— <code>vrf default</code> displays routing table for default VRF.</li> </ul> </li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1632.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1402; Arista User Manual, v. 4.11.1 (1/11/13), at 1148; Arista User Manual v. 4.10.3 (10/22/12), at 959.</p>	Dkt. 419-10 at PDF p. 343
<code>addr</code>	IPv4 address. The format is x.x.x.x																									
<code>advertised-routes</code>	(Optional) Displays all the routes advertised to this neighbor.																									
<code>flap-statistics</code>	(Optional) Displays flap statistics for the routes received from this neighbor.																									
<code>paths</code>	(Optional) Displays AS paths learned from this neighbor.																									
<code>received-routes</code>	(Optional) Displays all the routes received from this neighbor.																									
<code>routes</code>	(Optional) Displays the routes received or advertised to or from this neighbor.																									
<code>advertised</code>	(Optional) Displays all the routes advertised for this neighbor.																									
<code>damped</code>	(Optional) Displays all dampened routes received from this neighbor.																									
<code>received</code>	(Optional) Displays all the routes received from this neighbor.																									
<code>prefix</code>	(Optional) IPv6 prefix. The format is x.x.x.x/length.																									
<code>vrf vrf-name</code>	(Optional) Specifies the virtual router context (VRF) name. The name can be any case-sensitive, alphanumeric string up to 63 characters.																									
<code>all</code>	(Optional) Specifies all VRF.																									

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
<p>Use the <code>ip ospf database</code> command to display information about different OSPF LSAs.</p> <p>When the link state advertisement is describing a network, the <code>link-state-id</code> argument can take one of two forms:</p> <ul style="list-style-type: none"> <li>The network's IP address (such as Type 3 summary link advertisements and autonomous system external link advertisements).</li> <li>A derived address obtained from the link state ID. (Note that masking a network links advertisement's link state ID with the network's subnet mask yields the network's IP address.)</li> <li>When the link state advertisement is describing a router, the link state ID is always the described router's OSPF router ID.</li> <li>When an autonomous system external advertisement (LS Type = 5) is describing a default route, its link state ID is set to Default Destination (0.0.0.0).</li> </ul> <p>This command requires the Enterprise Services license.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 520.</p>	<ul style="list-style-type: none"> <li><i>linkstate_id</i> Network segment described by the LSA (dotted decimal notation). Value depends on the LSA type. <ul style="list-style-type: none"> <li>When the LSA describes a network, the <code>linkstate-id</code> argument is one of the following: <ul style="list-style-type: none"> <li>The network IP address, as in Type 3 summary link advertisements and in autonomous system external link advertisements.</li> <li>A derived address obtained from the link state ID. Masking a network links the advertisement link state ID with the network subnet mask yielding the network IP address.</li> </ul> </li> <li>When the LSA describes a router, the link state ID is the OSPFv2 router ID of the router.</li> <li>When an autonomous system external advertisement (Type 5) describes a default route, its link state ID is set to the default destination (0.0.0.0).</li> </ul> </li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1454.</p> <p><i>See also</i> Arista User Manual v. 4.13.6F (4/14/2014), at 1404; Arista User Manual v. 4.12.3 (7/17/13), at 1240; Arista User Manual, v. 4.11.1 (1/11/13), at 996; Arista User Manual v. 4.10.3 (10/22/12), at 825; Arista User Manual v. 4.9.3.2 (5/3/12), at 648; Arista User Manual v. 4.8.2 at 483; Arista User Manual v. 4.7.3 (7/18/11), at 357; Arista User Manual v. 4.6.0 (12/22/2010), at 217</p>	Dkt. 419-10 at PDF p. 344

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
<p>Use the <code>ip ospf database</code> command to display information about different OSPF LSAs.</p> <p>When the link state advertisement is describing a network, the <code>link-state-id</code> argument can take one of two forms:</p> <ul style="list-style-type: none"> <li>• The network's IP address (such as Type 3 summary link advertisements and autonomous system external link advertisements).</li> <li>• A derived address obtained from the link state ID. (Note that masking a network links advertisement's link state ID with the network's subnet mask yields the network's IP address.)</li> <li>• When the link state advertisement is describing a router, the link state ID is always the described router's OSPF router ID.</li> <li>• When an autonomous system external advertisement (LS Type = 5) is describing a default route, its link state ID is set to Default Destination (0.0.0.0).</li> </ul> <p>This command requires the Enterprise Services license.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 5.x (2010), at L3-742.</p>	<ul style="list-style-type: none"> <li>• <code>linkstate_id</code> Network segment described by the LSA (dotted decimal notation). Value depends on the LSA type. <ul style="list-style-type: none"> <li>— When the LSA describes a network, the <code>linkstate-id</code> argument is one of the following: <ul style="list-style-type: none"> <li>The network IP address, as in Type 3 summary link advertisements and in autonomous system external link advertisements.</li> <li>A derived address obtained from the link state ID. Masking a network links the advertisement link state ID with the network subnet mask yielding the network IP address.</li> </ul> </li> <li>— When the LSA describes a router, the link state ID is the OSPFv2 router ID of the router.</li> <li>— When an autonomous system external advertisement (Type 5) describes a default route, its link state ID is set to the default destination (0.0.0.0).</li> </ul> </li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1454.</p> <p><i>See also</i> Arista User Manual v. 4.13.6F (4/14/2014), at 1404; Arista User Manual v. 4.12.3 (7/17/13), at 1240; Arista User Manual, v. 4.11.1 (1/11/13), at 996; Arista User Manual v. 4.10.3 (10/22/12), at 825; Arista User Manual v. 4.9.3.2 (5/3/12), at 648; Arista User Manual v. 4.8.2 at 483; Arista User Manual v. 4.7.3 (7/18/11), at 357; Arista User Manual v. 4.6.0 (12/22/2010), at 217</p>	Dkt. 419-10 at PDF p. 345

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
<p>Use the <code>ip ospf database</code> command to display information about different OSPF LSAs.</p> <p>When the link state advertisement is describing a network, the <code>link-state-id</code> argument can take one of two forms:</p> <ul style="list-style-type: none"> <li>The network's IP address (such as Type 3 summary link advertisements and autonomous system external link advertisements).</li> <li>A derived address obtained from the link state ID. (Note that masking a network links advertisement's link state ID with the network's subnet mask yields the network's IP address.)</li> <li>When the link state advertisement is describing a router, the link state ID is always the described router's OSPF router ID.</li> <li>When an autonomous system external advertisement (LS Type = 5) is describing a default route, its link state ID is set to Default Destination (0.0.0.0).</li> </ul> <p>This command requires the Enterprise Services license.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 4.x (2008), at L3-426.</p>	<ul style="list-style-type: none"> <li><i>linkstate_id</i> Network segment described by the LSA (dotted decimal notation). Value depends on the LSA type. <ul style="list-style-type: none"> <li>When the LSA describes a network, the <code>linkstate-id</code> argument is one of the following: <ul style="list-style-type: none"> <li>The network IP address, as in Type 3 summary link advertisements and in autonomous system external link advertisements.</li> <li>A derived address obtained from the link state ID. Masking a network links the advertisement link state ID with the network subnet mask yielding the network IP address.</li> </ul> </li> <li>When the LSA describes a router, the link state ID is the OSPFv2 router ID of the router.</li> <li>When an autonomous system external advertisement (Type 5) describes a default route, its link state ID is set to the default destination (0.0.0.0).</li> </ul> </li> </ul> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1454.</p> <p><i>See also</i> Arista User Manual v. 4.13.6F (4/14/2014), at 1404; Arista User Manual v. 4.12.3 (7/17/13), at 1240; Arista User Manual, v. 4.11.1 (1/11/13), at 996; Arista User Manual v. 4.10.3 (10/22/12), at 825; Arista User Manual v. 4.9.3.2 (5/3/12), at 648; Arista User Manual v. 4.8.2 at 483; Arista User Manual v. 4.7.3 (7/18/11), at 357; Arista User Manual v. 4.6.0 (12/22/2010), at 217</p>	Dkt. 419-10 at PDF p. 346

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record									
<p><b>timers lsa-arrival (OSPF)</b></p> <p>To set the minimum interval in which the software accepts the same link-state advertisement (LSA) from Open Shortest Path First (OSPF) neighbors, use the <b>timers lsa-arrival</b> command. To return to the default, use the <b>no</b> form of this command.</p> <pre>timers lsa-arrival milliseconds no timers lsa-arrival</pre> <p><b>Syntax Description</b></p> <table border="1"> <tr> <td><i>milliseconds</i></td> <td>Minimum delay (in milliseconds) that must pass between acceptance of the same LSA arriving from neighbors. The range is from 10 to 600,000 milliseconds. The default is 1000 milliseconds.</td> </tr> </table> <p><b>Defaults</b></p> <table border="1"> <tr> <td>1000 milliseconds</td> </tr> </table> <p><b>Command Modes</b></p> <table border="1"> <tr> <td>Router configuration VRF configuration</td> </tr> </table> <p><b>SupportedUserRoles</b></p> <table border="1"> <tr> <td>network-admin vdc-admin</td> </tr> </table> <p><b>Command History</b></p> <table border="1"> <tr> <td><b>Release</b></td> <td><b>Modification</b></td> </tr> <tr> <td>4.0(1)</td> <td>This command was introduced.</td> </tr> </table> <p><b>Usage Guidelines</b></p> <p>Use the <b>timers lsa arrival</b> command to configure the minimum interval for accepting the same LSA. The same LSA is an LSA instance that contains the same LSA ID number, LSA type, and advertising router ID. If an instance of the same LSA arrives sooner than the interval that is set, the software drops the LSA.</p> <p>We recommend that you keep the <i>milliseconds</i> value of the <b>timers lsa-arrival</b> command less than or equal to the neighbors' <i>hold-interval</i> value of the <b>timers throttle lsa</b> command.</p> <p>This command requires the Enterprise Services license.</p> <p><b>Examples</b></p> <p>This example shows how to set the minimum interval for accepting the same LSA at 2000 milliseconds:</p> <pre>switch(config)# router ospf 1 switch(config-router)# timers lsa-arrival 2000</pre>	<i>milliseconds</i>	Minimum delay (in milliseconds) that must pass between acceptance of the same LSA arriving from neighbors. The range is from 10 to 600,000 milliseconds. The default is 1000 milliseconds.	1000 milliseconds	Router configuration VRF configuration	network-admin vdc-admin	<b>Release</b>	<b>Modification</b>	4.0(1)	This command was introduced.	<p><b>timers lsa arrival (OSPFv2)</b></p> <p>The <b>timers lsa arrival</b> command sets the minimum interval in which the switch accepts the same link-state advertisement (LSA) from OSPF neighbors.</p> <p>The no <b>timers lsa arrival</b> and <b>default timers lsa arrival</b> commands restore the default maximum OSPFv2 path calculation interval to five seconds by removing the <b>timers lsa arrival</b> command from <i>running-config</i>.</p> <p><b>Platform</b> all  <b>Command Mode</b> Router-OSPF Configuration</p> <p><b>Command Syntax</b></p> <pre>timers lsa arrival lsa_time no timers lsa arrival default timers lsa arrival</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li><i>lsa_time</i> OSPFv2 minimum interval (seconds). Values range from 1 to 600000 milliseconds. Default is 1000 milliseconds.</li> </ul> <p><b>Example</b></p> <ul style="list-style-type: none"> <li>This command sets the minimum interval timer to ten milliseconds.</li> </ul> <pre>switch(config)#router ospf 6 switch(config-router)#timers lsa arrival 10 switch(config-router-ospf)# </pre> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1469.</p>	Dkt. 419-10 at PDF p. 347
<i>milliseconds</i>	Minimum delay (in milliseconds) that must pass between acceptance of the same LSA arriving from neighbors. The range is from 10 to 600,000 milliseconds. The default is 1000 milliseconds.										
1000 milliseconds											
Router configuration VRF configuration											
network-admin vdc-admin											
<b>Release</b>	<b>Modification</b>										
4.0(1)	This command was introduced.										

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record									
<p><b>timers lsa-arrival (OSPF)</b></p> <p>To set the minimum interval in which the software accepts the same link-state advertisement (LSA) from Open Shortest Path First (OSPF) neighbors, use the <b>timers lsa-arrival</b> command. To return to the default, use the <b>no</b> form of this command.</p> <pre>timers lsa-arrival milliseconds no timers lsa-arrival</pre> <p><b>Syntax Description</b></p> <table border="1"> <tr> <td><i>milliseconds</i></td> <td>Minimum delay (in milliseconds) that must pass between acceptance of the same LSA arriving from neighbors. The range is from 10 to 600,000 milliseconds. The default is 1000 milliseconds.</td> </tr> </table> <p><b>Defaults</b></p> <table border="1"> <tr> <td>1000 milliseconds</td> </tr> </table> <p><b>Command Modes</b></p> <table border="1"> <tr> <td>Router configuration VRF configuration</td> </tr> </table> <p><b>SupportedUserRoles</b></p> <table border="1"> <tr> <td>network-admin vdc-admin</td> </tr> </table> <p><b>Command History</b></p> <table border="1"> <tr> <td><b>Release</b></td> <td><b>Modification</b></td> </tr> <tr> <td>4.0(1)</td> <td>This command was introduced.</td> </tr> </table> <p><b>Usage Guidelines</b></p> <p>Use the <b>timers lsa arrival</b> command to configure the minimum interval for accepting the same LSA. The same LSA is an LSA instance that contains the same LSA ID number, LSA type, and advertising router ID. If an instance of the same LSA arrives sooner than the interval that is set, the software drops the LSA.</p> <p>We recommend that you keep the <i>milliseconds</i> value of the <b>timers lsa-arrival</b> command less than or equal to the neighbors' <i>hold-interval</i> value of the <b>timers throttle lsa</b> command.</p> <p>This command requires the Enterprise Services license.</p> <p><b>Examples</b></p> <p>This example shows how to set the minimum interval for accepting the same LSA at 2000 milliseconds:</p> <pre>switch(config)# router ospf 1 switch(config-router)# timers lsa-arrival 2000</pre>	<i>milliseconds</i>	Minimum delay (in milliseconds) that must pass between acceptance of the same LSA arriving from neighbors. The range is from 10 to 600,000 milliseconds. The default is 1000 milliseconds.	1000 milliseconds	Router configuration VRF configuration	network-admin vdc-admin	<b>Release</b>	<b>Modification</b>	4.0(1)	This command was introduced.	<p><b>timers lsa arrival (OSPFv2)</b></p> <p>The <b>timers lsa arrival</b> command sets the minimum interval in which the switch accepts the same link-state advertisement (LSA) from OSPF neighbors.</p> <p>The no <b>timers lsa arrival</b> and <b>default timers lsa arrival</b> commands restore the default maximum OSPFv2 path calculation interval to five seconds by removing the <b>timers lsa arrival</b> command from <i>running-config</i>.</p> <p><b>Platform</b> all  <b>Command Mode</b> Router-OSPF Configuration</p> <p><b>Command Syntax</b></p> <pre>timers lsa arrival lsa_time no timers lsa arrival default timers lsa arrival</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li><i>lsa_time</i> OSPFv2 minimum interval (seconds). Values range from 1 to 600000 milliseconds. Default is 1000 milliseconds.</li> </ul> <p><b>Example</b></p> <ul style="list-style-type: none"> <li>This command sets the minimum interval timer to ten milliseconds.</li> </ul> <pre>switch(config)#router ospf 6 switch(config-router)#timers lsa arrival 10 switch(config-router-ospf)# </pre> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1469.</p>	Dkt. 419-10 at PDF p. 348
<i>milliseconds</i>	Minimum delay (in milliseconds) that must pass between acceptance of the same LSA arriving from neighbors. The range is from 10 to 600,000 milliseconds. The default is 1000 milliseconds.										
1000 milliseconds											
Router configuration VRF configuration											
network-admin vdc-admin											
<b>Release</b>	<b>Modification</b>										
4.0(1)	This command was introduced.										

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record									
<p><b>timers lsa-arrival (OSPF)</b></p> <p>To set the minimum interval in which the software accepts the same link-state advertisement (LSA) from Open Shortest Path First (OSPF) neighbors, use the <b>timers lsa-arrival</b> command. To return to the default, use the <b>no</b> form of this command.</p> <pre>timers lsa-arrival milliseconds no timers lsa-arrival</pre> <p><b>Syntax Description</b></p> <table border="1"> <tr> <td><i>milliseconds</i></td> <td>Minimum delay (in milliseconds) that must pass between acceptance of the same LSA arriving from neighbors. The range is from 10 to 600,000 milliseconds. The default is 1000 milliseconds.</td> </tr> </table> <p><b>Defaults</b></p> <table border="1"> <tr> <td>1000 milliseconds</td> </tr> </table> <p><b>Command Modes</b></p> <table border="1"> <tr> <td>Router configuration VRF configuration</td> </tr> </table> <p><b>SupportedUserRoles</b></p> <table border="1"> <tr> <td>network-admin vdc-admin</td> </tr> </table> <p><b>Command History</b></p> <table border="1"> <tr> <td><b>Release</b></td> <td><b>Modification</b></td> </tr> <tr> <td>4.0(1)</td> <td>This command was introduced.</td> </tr> </table> <p><b>Usage Guidelines</b></p> <p>Use the <b>timers lsa arrival</b> command to configure the minimum interval for accepting the same LSA. The same LSA is an LSA instance that contains the same LSA ID number, LSA type, and advertising router ID. If an instance of the same LSA arrives sooner than the interval that is set, the software drops the LSA.</p> <p>We recommend that you keep the <i>milliseconds</i> value of the <b>timers lsa-arrival</b> command less than or equal to the neighbors' <i>hold-interval</i> value of the <b>timers throttle lsa</b> command.</p> <p>This command requires the Enterprise Services license.</p> <p><b>Examples</b></p> <p>This example shows how to set the minimum interval for accepting the same LSA at 2000 milliseconds:</p> <pre>switch(config)# router ospf 1 switch(config-router)# timers lsa-arrival 2000</pre>	<i>milliseconds</i>	Minimum delay (in milliseconds) that must pass between acceptance of the same LSA arriving from neighbors. The range is from 10 to 600,000 milliseconds. The default is 1000 milliseconds.	1000 milliseconds	Router configuration VRF configuration	network-admin vdc-admin	<b>Release</b>	<b>Modification</b>	4.0(1)	This command was introduced.	<p><b>timers lsa arrival (OSPFv2)</b></p> <p>The <b>timers lsa arrival</b> command sets the minimum interval in which the switch accepts the same link-state advertisement (LSA) from OSPF neighbors.</p> <p>The no <b>timers lsa arrival</b> and <b>default timers lsa arrival</b> commands restore the default maximum OSPFv2 path calculation interval to five seconds by removing the <b>timers lsa arrival</b> command from <i>running-config</i>.</p> <p><b>Platform</b> all  <b>Command Mode</b> Router-OSPF Configuration</p> <p><b>Command Syntax</b></p> <pre>timers lsa arrival lsa_time no timers lsa arrival default timers lsa arrival</pre> <p><b>Parameters</b></p> <ul style="list-style-type: none"> <li><i>lsa_time</i> OSPFv2 minimum interval (seconds). Values range from 1 to 600000 milliseconds. Default is 1000 milliseconds.</li> </ul> <p><b>Example</b></p> <ul style="list-style-type: none"> <li>This command sets the minimum interval timer to ten milliseconds.</li> </ul> <pre>switch(config)#router ospf 6 switch(config-router)#timers lsa arrival 10 switch(config-router-ospf)# </pre> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1469.</p>	Dkt. 419-10 at PDF p. 349
<i>milliseconds</i>	Minimum delay (in milliseconds) that must pass between acceptance of the same LSA arriving from neighbors. The range is from 10 to 600,000 milliseconds. The default is 1000 milliseconds.										
1000 milliseconds											
Router configuration VRF configuration											
network-admin vdc-admin											
<b>Release</b>	<b>Modification</b>										
4.0(1)	This command was introduced.										

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
<p><b>Examples</b></p> <p>This example shows how to configure a router configured with the start, hold, and maximum interval values for the <b>timers throttle spf</b> command set at 5, 1000, and 90,000 milliseconds:</p> <pre>switch(config)# router ospf 1 switch(config-router)# timers throttle spf 5 1000 90000</pre> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 1033-34.</p> <p>When using route reflectors, an AS is divided into clusters. A cluster consists of one or more route reflectors and a group of clients to which they re-advertise route information. Multiple route reflectors can be configured in the same cluster to increase redundancy and avoid a single point of failure. Each route reflector has a cluster ID. If the cluster has a single route reflector, the cluster ID is its router ID. If a cluster has multiple route reflectors, a 4-byte cluster ID is assigned to all route reflectors in the cluster. All of them must be configured with the same cluster ID so that they can recognize updates from other route reflectors in the same cluster. The <b>bgp cluster-id</b> command configures the cluster ID in a cluster with multiple route reflectors.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 730.</p>	<p><b>Example</b></p> <ul style="list-style-type: none"> <li>This command displays a switch configured with the start, hold, and maximum interval values for the timers throttle spf command set at 5, 1,000, and 20,000 milliseconds, respectively.</li> </ul> <pre>switch(config)#router ospf 6 switch(config-router)#timers spf 5 100 20000 switch(config-router-ospf)# </pre> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1472.</p>	Dkt. 419-10 at PDF p. 350
<p><b>Local Proxy ARP</b></p> <p>You can use local Proxy ARP to enable a device to respond to ARP requests for IP addresses within a subnet where normally no routing is required. When you enable local Proxy ARP, ARP responds to all ARP requests for IP addresses within the subnet and forwards all traffic between hosts in the subnet. Use this feature only on subnets where hosts are intentionally prevented from communicating directly by the configuration on the device to which they are connected.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x (2013), at 2-5.</p>	<p><b>ip local-proxy-arp</b></p> <p>The <b>ip local-proxy-arp</b> command enables local proxy ARP (Address Resolution Protocol) on the configuration mode interface. Local proxy ARP programs the switch to respond to ARP requests for IP addresses within a subnet where routing is not normally required. A typical local proxy arp application is supporting isolated private VLANs that communicate with each other by routing packets.</p> <p>The <b>no ip local-proxy-arp</b> and <b>default ip local-proxy-arp</b> commands disable local proxy ARP on the configuration mode interface by removing the corresponding <b>ip local-proxy-arp</b> command from <b>running-config</b>.</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1276.</p> <p><i>See also</i> Arista User Manual v. 4.13.6F (4/14/2014), at 1231; Arista User Manual v. 4.12.3 (7/17/13), at 1073; Arista User Manual, v. 4.11.1 (1/11/13), at 876; Arista User Manual v. 4.10.3 (10/22/12), at 707.</p>	Dkt. 419-10 at PDF p. 350

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
<p><b>Local Proxy ARP</b></p> <p>You can use local Proxy ARP to enable a device to respond to ARP requests for IP addresses within a subnet where normally no routing is required. When you enable local Proxy ARP, ARP responds to all ARP requests for IP addresses within the subnet and forwards all traffic between hosts in the subnet. Use this feature only on subnets where hosts are intentionally prevented from communicating directly by the configuration on the device to which they are connected.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x (2010), at 2-5.</p>	<p><b>ip local-proxy-arp</b></p> <p>The ip local-proxy-arp command enables local proxy ARP (Address Resolution Protocol) on the configuration mode interface. Local proxy ARP programs the switch to respond to ARP requests for IP addresses within a subnet where routing is not normally required. A typical local proxy arp application is supporting isolated private VLANs that communicate with each other by routing packets.</p> <p>The no ip local-proxy-arp and default ip local-proxy-arp commands disable local proxy ARP on the configuration mode interface by removing the corresponding ip local-proxy-arp command from running-config.</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1276.</p> <p><i>See also</i> Arista User Manual v. 4.13.6F (4/14/2014), at 1231; Arista User Manual v. 4.12.3 (7/17/13), at 1073; Arista User Manual, v. 4.11.1 (1/11/13), at 876; Arista User Manual v. 4.10.3 (10/22/12), at 707.</p>	Dkt. 419-10 at PDF p. 351
<p><b>Local Proxy ARP</b></p> <p>You can use local Proxy ARP to enable a device to respond to ARP requests for IP addresses within a subnet where normally no routing is required. When you enable local Proxy ARP, ARP responds to all ARP requests for IP addresses within the subnet and forwards all traffic between hosts in the subnet. Use this feature only on subnets where hosts are intentionally prevented from communicating directly by the configuration on the device to which they are connected.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 4.0 (2008), at 2-5.</p>	<p><b>ip local-proxy-arp</b></p> <p>The ip local-proxy-arp command enables local proxy ARP (Address Resolution Protocol) on the configuration mode interface. Local proxy ARP programs the switch to respond to ARP requests for IP addresses within a subnet where routing is not normally required. A typical local proxy arp application is supporting isolated private VLANs that communicate with each other by routing packets.</p> <p>The no ip local-proxy-arp and default ip local-proxy-arp commands disable local proxy ARP on the configuration mode interface by removing the corresponding ip local-proxy-arp command from running-config.</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1276.</p> <p><i>See also</i> Arista User Manual v. 4.13.6F (4/14/2014), at 1231; Arista User Manual v. 4.12.3 (7/17/13), at 1073; Arista User Manual, v. 4.11.1 (1/11/13), at 876; Arista User Manual v. 4.10.3 (10/22/12), at 707.</p>	Dkt. 419-10 at PDF p. 351

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
<p><code>ipv6 nd managed-config-flag</code> Sets the managed address configuration flag in IPv6 router advertisements.</p> <p><code>ipv6 nd mtu</code> Sets the maximum transmission unit (MTU) size of IPv6 packets sent on an interface.</p> <p><code>ipv6 nd ns-interval</code> Configures the interval between IPv6 neighbor solicitation retransmissions on an interface.</p> <p><code>ipv6 nd other-config-flag</code> Configures the other stateful configuration flag in IPv6 router advertisements.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x (2013), at 3-24.</p>	<p><b>Router Advertisement Flag Configuration</b></p> <p>The <code>ipv6 nd managed-config-flag</code> command configures the switch to set the <i>managed address configuration</i> flag in IPv6 router advertisements transmitted from the configuration mode interface. This bit instructs receptive hosts to use stateful address autoconfiguration.</p> <p>The <code>ipv6 nd other-config-flag</code> command configures the switch to set the <i>other stateful configuration</i> flag in IPv6 router advertisements transmitted from the configuration mode interface. This flag indicates the availability of autoconfiguration information, other than addresses, and that hosts should use stateful</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1329.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1119; Arista User Manual, v. 4.11.1 (1/11/13), at 887; Arista User Manual v. 4.10.3 (10/22/12), at 733.</p>	Dkt. 419-10 at PDF p. 352
<p><code>ipv6 nd managed-config-flag</code> Sets the managed address configuration flag in IPv6 router advertisements.</p> <p><code>ipv6 nd mtu</code> Sets the maximum transmission unit (MTU) size of IPv6 packets sent on an interface.</p> <p><code>ipv6 nd ns-interval</code> Configures the interval between IPv6 neighbor solicitation retransmissions on an interface.</p> <p><code>ipv6 nd other-config-flag</code> Configures the other stateful configuration flag in IPv6 router advertisements.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x (2010), at 3-22.</p>	<p><b>Router Advertisement Flag Configuration</b></p> <p>The <code>ipv6 nd managed-config-flag</code> command configures the switch to set the <i>managed address configuration</i> flag in IPv6 router advertisements transmitted from the configuration mode interface. This bit instructs receptive hosts to use stateful address autoconfiguration.</p> <p>The <code>ipv6 nd other-config-flag</code> command configures the switch to set the <i>other stateful configuration</i> flag in IPv6 router advertisements transmitted from the configuration mode interface. This flag indicates the availability of autoconfiguration information, other than addresses, and that hosts should use stateful</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1329.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1119; Arista User Manual, v. 4.11.1 (1/11/13), at 887; Arista User Manual v. 4.10.3 (10/22/12), at 733.</p>	Dkt. 419-10 at PDF p. 352

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
<p><code>ipv6 nd managed-config-flag</code> Sets the managed address configuration flag in IPv6 router advertisements.</p> <p><code>ipv6 nd mtu</code> Sets the maximum transmission unit (MTU) size of IPv6 packets sent on an interface.</p> <p><code>ipv6 nd ns-interval</code> Configures the interval between IPv6 neighbor solicitation retransmissions on an interface.</p> <p><code>ipv6 nd other-config-flag</code> Configures the other stateful configuration flag in IPv6 router advertisements.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 4.0 (2008), at 3-22.</p>	<p><b>Router Advertisement Flag Configuration</b></p> <p>The <code>ipv6 nd managed-config-flag</code> command configures the switch to set the <i>managed address configuration</i> flag in IPv6 router advertisements transmitted from the configuration mode interface. This bit instructs receptive hosts to use stateful address autoconfiguration.</p> <p>The <code>ipv6 nd other-config-flag</code> command configures the switch to set the <i>other stateful configuration flag</i> in IPv6 router advertisements transmitted from the configuration mode interface. This flag indicates the availability of autoconfiguration information, other than addresses, and that hosts should use stateful</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1329.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1119; Arista User Manual, v. 4.11.1 (1/11/13), at 887; Arista User Manual v. 4.10.3 (10/22/12), at 733.</p>	Dkt. 419-10 at PDF p. 353
<p><code>ipv6 nd reachable-time</code> Configures the amount of time that a remote IPv6 node is considered reachable after some reachability confirmation event has occurred.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x (2013), at 3-24.</p>	<p><b>ipv6 nd reachable-time</b></p> <p>The <code>ipv6 nd reachable-time</code> command specifies the time period that the switch includes in the reachable time field of Router Advertisements (RAs) sent from the configuration mode interface. The reachable time defines the period that a remote IPv6 node is considered reachable after a reachability confirmation event.</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1359.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1149.</p>	Dkt. 419-10 at PDF p. 353
<p><code>ipv6 nd reachable-time</code> Configures the amount of time that a remote IPv6 node is considered reachable after some reachability confirmation event has occurred.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x (2010), at 3-22.</p>	<p><b>ipv6 nd reachable-time</b></p> <p>The <code>ipv6 nd reachable-time</code> command specifies the time period that the switch includes in the reachable time field of Router Advertisements (RAs) sent from the configuration mode interface. The reachable time defines the period that a remote IPv6 node is considered reachable after a reachability confirmation event.</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1359.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1149.</p>	Dkt. 419-10 at PDF p. 353

Cisco's Documentation	Arista's Documentation	Supporting Evidence In The Record
<p><b>ipv6 nd reachable-time</b></p> <p>Configures the amount of time that a remote IPv6 node is considered reachable after some reachability confirmation event has occurred.</p> <p>Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 4.0 (2008), at 3-22.</p>	<p><b>ipv6 nd reachable-time</b></p> <p>The <b>ipv6 nd reachable-time</b> command specifies the time period that the switch includes in the reachable time field of Router Advertisements (RAs) sent from the configuration mode interface. The reachable time defines the period that a remote IPv6 node is considered reachable after a reachability confirmation event.</p> <p>Arista User Manual v. 4.14.3F – Rev. 2 (10/2/2014), at 1359.</p> <p><i>See also</i> Arista User Manual v. 4.12.3 (7/17/13), at 1149.</p>	Dkt. 419-10 at PDF p. 354